

APPEAL NO.
2015-1719

UNITED STATES COURT OF APPEALS FOR THE FEDERAL CIRCUIT

DELL INC., HEWLETT-PACKARD COMPANY, AND
NETAPP, INC.
Appellants,
v.

ELECTRONICS AND TELECOMMUNICATIONS RESEARCH INSTITUTE,
Appellee.

On appeal from the United States Patent and Trademark Office
Patent Trial and Appeal Board in IPR2013-00635

BRIEF FOR APPELLANTS
DELL INC., HEWLETT-PACKARD COMPANY, AND NETAPP, INC.

Thomas M. Dunham
James M. Woods
Winston & Strawn LLP
1700 K Street, NW
Washington, DC 20006
Telephone: (202) 282-5863

Counsel for Appellant Dell Inc.

Peter H. Kang
SIDLEY AUSTIN LLP
1001 Page Mill Road, Building 1
Palo Alto, CA 94304
Telephone: (650) 565-7006

Anna M. Weinberg
SIDLEY AUSTIN LLP
1501 K Street, N.W.
Washington, D.C. 20005
Telephone: (202) 736-8000

*Counsel for Appellants Dell Inc.,
Hewlett Packard Enterprise Company,
and NetApp, Inc.*

CERTIFICATE OF INTEREST

UNITED STATES COURT OF APPEALS FOR THE FEDERAL CIRCUIT

Dell Inc., Hewlett-Packard Company, and NetApp, Inc. v. Electronics and Telecommunications Research Institute

No. 15-1719

CERTIFICATE OF INTEREST

Counsel for the (petitioner) (appellant) (respondent) (appellee) (amicus) (name of party)

Dell Inc., Hewlett-Packard Company, and NetApp, Inc. certifies the following (use "None" if applicable; use extra sheets if necessary):

1. The full name of every party or amicus represented by me is:

Dell Inc., Hewlett Packard Enterprise Company (f/k/a Hewlett-Packard Company), and NetApp, Inc.

2. The name of the real party in interest (if the party named in the caption is not the real party in interest) represented by me is:

Dell Inc. is a privately held corporation and its direct parent company is Denali Intermediate Inc.

3. All parent corporations and any publicly held companies that own 10 percent or more of the stock of the party or amicus curiae represented by me are:

As of October 29, 2013, Dell Inc. is a privately held corporation and its direct parent company is Denali Intermediate Inc. There is no publicly held company owning 10% or more of Denali Intermediate Inc.'s stock. Ownership documentation is still being completed for Hewlett Packard Enterprise Company ("HPE"), and upon completion, Appellants will file an updated Certificate of Interest reflecting any entities that own more than 10% of HPE.

4. ☒ The names of all law firms and the partners or associates that appeared for the party or amicus now represented by me in the trial court or agency or are expected to appear in this court are:

Haynes Boone, LLP - David McCombs, Andrew S. Ehmke, John R. Emerson and Thomas W. Kelton; Morgan Lewis & Bockius LLP - Jody Barillare, Mark W. Taylor, Michael Lyons, Walter S. Tester and John Gorman; Sidley Austin LLP - Peter H. Kang and Anna M. Weinberg; Winston & Strawn LLP - Thomas M. Dunham and J. Michael Woods

11/11/2015

Date

/s/ Peter H. Kang

Signature of counsel

Peter H. Kang

Printed name of counsel

Please Note: All questions must be answered

cc: See attached Certificate of Service

TABLE OF CONTENTS

STATEMENT OF RELATED CASES	xi
JURISDICTIONAL STATEMENT	1
STATEMENT OF ISSUES	1
INTRODUCTION	2
STATEMENT OF THE CASE.....	4
I. Statement of Facts.	4
A. The '346 Patent	4
1. The '346 Patent Specification Describes Several Prior Art Embodiments Utilizing RAID.	5
2. The '346 Patent Specification Describes Providing Fault Tolerance of RAID Controllers Without Sacrificing Performance.	9
3. The '346 Patent Claims at Issue.	11
B. Proceedings Below	12
1. Petition for <i>Inter Partes</i> Review	12
2. The Board's Institution Based on Hathorn and the Institution Decision's Construction of "RAID".	13
3. Disputes Over Construction of "RAID" Following Institution.	15
4. The Board's Final Written Decision Narrows the Construction of "RAID" and Reverses the Initial Decision That Hathorn Likely Anticipates.	18
a. The Board's Final Written Decision Adopted a New, Narrow Construction of "RAID" Based on Select, Partial Dictionary Definitions.	18
b. The Board's Final Written Decision Relied on This New Narrowed Definition of "RAID" to Conclude that Hathorn Does Not Anticipate the '346 Claims.	20
C. Hathorn	21
D. Extrinsic Evidence from the Record Below	24
1. Prior Art References.....	24

a.	Weygant.....	24
b.	Chen.....	26
2.	Dictionaries	29
a.	Microsoft Computer Dictionary (1999).....	29
b.	Webster’s Computer Dictionary (2001)	30
c.	Microsoft Computer Dictionary (2002).....	30
II.	SUMMARY OF ARGUMENT.....	31
III.	ARGUMENT.....	34
A.	Standard of Review.	34
B.	The Board’s Narrowed Construction of “RAID” is Inconsistent with the ’346 Patent and How a Person of Ordinary Skill in the Art Would Understand the Term.....	35
1.	The Broadest Reasonable Interpretation of “RAID” Is “Redundant Array of Inexpensive Disks”	37
a.	The Intrinsic Evidence Supports Construing “RAID” to Mean a “Redundant Array of Inexpensive Disks”	37
b.	Extrinsic Evidence Confirms That the Plain And Ordinary Meaning of “RAID” Is a “Redundant Array of Inexpensive Disks”	44
c.	The Extrinsic Record Does Not Support the Board’s Construction.....	48
d.	Neither the Intrinsic Nor Extrinsic Record Supports the Board’s Further Narrowing of “RAID” to Preclude Direct Communications With Storage Controllers.....	51
C.	Hathorn Anticipates the ’346 Patent Challenged Claims Under Both the Proper Construction and the Board’s Erroneous Construction of “RAID”	52
1.	Hathorn Discloses a System with a Redundant Array of Inexpensive Disks In a Mirrored Configuration.....	52
2.	Hathorn Discloses a “RAID” Even Under the Board’s Erroneous Construction	56
IV.	CONCLUSION	61

TABLE OF AUTHORITIES

	Page(s)
 Cases	
<i>01 Communique Lab., Inc. v. LogMeIn, Inc.</i> , 687 F.3d 1292 (Fed. Cir. 2012), <i>aff'd</i> , 563 F. App'x 770 (Fed. Cir. 2014)	43
<i>3M Innovative Props. Co. v. Tredegar Corp.</i> , 725 F.3d 1315 (Fed. Cir. 2013)	33, 41, 42, 43
<i>Arthrocare Corp. v. Smith & Nephew, Inc.</i> , 406 F.3d 1365 (Fed. Cir. 2005)	60
<i>Baldwin Graphic Sys., Inc. v. Siebert, Inc.</i> , 512 F.3d 1338 (Fed. Cir. 2008)	43
<i>In re Baxter Travenol Labs.</i> , 952 F.2d 388 390 (Fed. Cir. 1991)	35
<i>Briantree Labs., Inc. v. Novel Labs., Inc.</i> , 749 F.3d 1349 (Fed. Cir. 2014)	56
<i>Combined Sys., Inc. v. Defense Tech. Corp. of Am.</i> , 350 F.3d 1207 (Fed. Cir. 2003)	48
<i>In re Cuozzo Speed Techs., LLC</i> , 793 F.3d 1268 (Fed. Cir. 2015)	34
<i>Helmsderfer v. Bobrick Washroom Equip., Inc.</i> , 527 F.3d 1379 (Fed. Cir. 2008)	44, 48
<i>Hewlett-Packard Co. v. Mustek Sys., Inc.</i> , 340 F.3d 1314 (Fed. Cir. 2003)	60
<i>InterDigital Comm'ns, LLC v. ITC</i> , 690 F.3d 1318 (Fed. Cir. 2012)	35, 38, 42, 50
<i>Interval Licensing LLC v. AOL, Inc.</i> , 766 F.3d 1364 (Fed. Cir. 2014)	48

<i>Mass. Inst. of Tech. & Elec. For Imaging, Inc. v. Abacus Software,</i> 462 F.3d 1344 (Fed. Cir. 2006)	46
<i>MBO Labs., Inc. v. Becton, Dickenson & Co.,</i> 474 F.3d 1323 (Fed. Cir. 2007)	37
<i>Nystrom v. TREX Co.,</i> 424 F.3d 1136 (Fed. Cir. 2005)	40
<i>Phillips v. AWH Corp.,</i> 415 F.3d 1303 (Fed. Cir. 2005)	37, 49
<i>Power Integrations, Inc. v. Fairchild Semiconductor Int’l, Inc.,</i> 711 F.3d 1348 (Fed. Cir. 2013)	48
<i>Southwall Techs., Inc. v. Cardinal IG Co.,</i> 54 F.3d 1570 (Fed. Cir. 1995)	40
<i>Starhome GmnH v. AT&T Mobility LLC,</i> 743 F.3d 849 (Fed. Cir. 2014)	46
<i>Streck, Inc. v. Research & Diagnostic Sys., Inc.,</i> 665 F.3d 1269 (Fed. Cir. 2012)	41
<i>In re Suitco Surface, Inc.,</i> 603 F.3d 1255 (Fed. Cir. 2010)	35
<i>Tegal Corp. v. Tokyo Electron Am., Inc.,</i> 275 F.3d 1331 (Fed. Cir. 2001)	33
<i>Teleflex, Inc. v. Ficosa N. Am. Corp.,</i> 299 F.3d 1313 (Fed. Cir. 2002)	38
<i>Tempo Lighting, Inc. v. Tivoli, LLC,</i> 742 F.3d 973 (Fed. Cir. 2014)	44
<i>Teva Pharm. USA, Inc. v. Sandoz, Inc.,</i> 135 S. Ct. 831 (2015)	34
<i>Texas Digital Sys., Inc. v. Telegenix, Inc.,</i> 308 F.3d 1193 (Fed. Cir. 2002)	45

<i>TI Grp. Auto. Sys. (N.A.), Inc. v. VDO N.A., LLC</i> , 375 F.3d 1126 (Fed. Cir. 2004)	56
<i>TiVo, Inc. v. EchoStar Commc'ns Corp.</i> , 516 F.3d 1290 (Fed. Cir. 2008)	43
<i>W.E. Hall Co. v. Atlanta Corrugating, LLC</i> , 370 F.3d 1343 (Fed. Cir. 2004)	44
Statutes	
28 U.S.C. § 1295(a)(4)(A)	1
35 U.S.C. § 6	1
35 U.S.C. § 318(a)	1
35 U.S.C. § 319	1
Other Authorities	
37 C.F.R. § 42.100(b)	34

STATEMENT OF RELATED CASES

The asserted claims of the patent at issue in this appeal, U.S. Patent No. 6,978,346 (“the ’346 patent”), are also the subject of two other *inter partes* review proceedings before the Patent and Trial Appeal Board (“Board”): IPR2014-00901, which was filed by VMware, Inc. on June 4, 2014, and which is currently pending Final Written Decision, and IPR2014-00949, which was filed by International Business Machines Corporation and Oracle America, Inc. on June 13, 2014 and which is currently pending Final Written Decision.

In addition to these *inter partes* review proceedings, several district court cases pending before the U.S. District Court for the District of Delaware involve the ’346 patent: *Safe Storage LLC v. Int’l Business Machine Corp.*, 1-13-cv-01151 (filed June 28, 2013); *Safe Storage LLC v. Emulex Corp. et al.*, 1-13-cv-01150 (filed June 28, 2013); *Safe Storage LLC v. ATTO Tech. Inc. et al.*, 1-13-cv-01090 (filed June 17, 2013); *Safe Storage LLC v. Oracle Am. Inc. et al.*, 1-13-cv-01089 (filed June 17, 2013); *Safe Storage LLC v. Overland Storage Inc.*, 1-13-cv-00932 (filed May 23, 2013); *Safe Storage LLC v. Infortrend Corp.*, 1-13-cv-00929 (filed May 23, 2013); *Safe Storage LLC v. Nexsan Corp.*, 1-13-cv-00931 (filed May 23, 2013); *Safe Storage LLC v. VMware Inc.*, 1-13-cv-00928 (filed May 23, 2013); *Safe Storage LLC v. Cisco Sys. Inc.*, 1-13-cv-00926 (filed May 23, 2013); *Safe Storage LLC v. Dell Inc.*, 1-12-cv-01624 (filed Nov. 30, 2012); *Safe Storage LLC*

v. Silicon Graphics Int’l Corp., 1-12-cv-01629 (filed Nov. 30, 2012); *Safe Storage LLC v. Hewlett-Packard Co.*, 1-12-cv-01626 (filed Nov. 30, 2012); *Safe Storage LLC v. Dot Hill Sys Corp.*, 1-12-cv-01625 (filed Nov. 30, 2012); *Safe Storage LLC v. Netapp Inc.*, 1-12-cv-01628 (filed Nov. 30, 2012); and *Safe Storage LLC v. Silicon Graphics Int’l Corp.*, 1-12-cv-01629 (filed Nov. 30, 2012). To the best of Appellants’ knowledge at the time of filing, all of the above mentioned district court litigations have been stayed pending disposition of the various *inter partes* review proceedings challenging the validity of the ’346 patent.

JURISDICTIONAL STATEMENT

The Patent Trial and Appeal Board (“Board”) had jurisdiction over this *inter partes* review proceeding under 35 U.S.C. § 6. On March 20, 2014, the Board instituted trial for claims 1-3 and 5-8 of the ’346 patent (the “Institution Decision”). The Board entered its Final Written Decision pursuant to 35 U.S.C. § 318(a) on February 27, 2015. Appellants Dell Inc., Hewlett Packard Enterprise Company,¹ and NetApp Inc. (“Appellants” or “Petitioners”) filed a timely notice of appeal on April 30, 2015, appealing to this Court the Board’s determination that claims 1-3 and 5-8 of the ’346 patent have not been shown to be unpatentable by a preponderance of the evidence. This Court has jurisdiction under 28 U.S.C. § 1295(a)(4)(A) and 35 U.S.C. § 319.

STATEMENT OF ISSUES

1. Under the broadest reasonable interpretation (“BRI”) standard, whether the Board’s Institution Decision properly interpreted the ’346 patent claim term “RAID” consistent with the ’346 patent specification to mean a “redundant array of inexpensive disks”?

¹ On November 1, 2015, Hewlett-Packard Company separated into two new companies. See Dkt. No. 28 (Nov. 13, 2015) (Unopposed Motion to Reform Official Caption). Hewlett Packard Enterprise Company is the successor to Hewlett-Packard Company and real party in interest for purposes of this appeal. See Dkt. No. 27 (Nov. 13, 2015) (Amended Certificate of Interest for the Appellants).

2. Under the BRI standard, whether the Board failed to give the claim term “RAID” a construction consistent with the specification when, in its Final Written Decision, the Board narrowed its construction of “RAID” to “a single logical unit for mass storage using multiple physical disk drives” and further excluded RAIDs with known “mirroring configurations” from the scope of this claim term, based on improperly importing limitations from extrinsic evidence and without regard for the broad disclosure of “RAID” in the intrinsic evidence?

3. Whether U.S. Patent No. 5,574,950 issued to Hathorn et al. (“Hathorn”), which indisputably qualifies as prior art to the ’346 patent, anticipates claims 1-3 and 5-8 of the ’346 patent, either under (a) the Board’s initial interpretation in the Institution, and Appellants’ construction, of “RAID” as a “redundant array of inexpensive disks” or (b) the Board’s improper extrinsic evidence-based construction of “RAID” as “a single logical unit for mass storage using multiple physical disk drives,” where no one disputes that Hathorn discloses a RAID system which, in one preferred embodiment (ignored by the Board), can be configured to act like a “single logical unit”?

INTRODUCTION

This appeal centers around the proper claim construction of the term “RAID.” The ’346 specification defines the well-known industry acronym “RAID” as a “redundant array of inexpensive disks,” which is the plain and

ordinary meaning of the term to a person of ordinary skill. (JA189). In the Institution Decision, the Board adopted this correct construction of “RAID.” (JA189). But in its Final Written Decision, the Board improperly narrowed its construction to “a single logical unit for mass storage using multiple physical disk drives” in contravention of the BRI standard (JA5-9). In doing so, the Board legally erred by importing limitations from the extrinsic record into the term to narrow its plain and ordinary meaning, particularly where such importation was inconsistent with the intrinsic evidence and is based on selective interpretations of a RAID in the extrinsic record.

Because the Board relied on this improper construction, this Court should also reverse the Board’s ruling on unpatentability. Specifically, the improper narrow construction of “RAID” led the Board to the conclusion that the Hathorn prior art reference does not anticipate the ’346 patent claims. But the Board failed to take into account all embodiments disclosed in Hathorn which, even under the Board’s incorrect construction, still satisfy the claim term “RAID.” And under the proper construction of “RAID,” Hathorn anticipates the claims at issue.

Thus, the Court should reverse the Board’s Final Written Decision both as to the claim construction of “RAID” and as to anticipation by Hathorn.

STATEMENT OF THE CASE

I. Statement of Facts.

A. The '346 Patent.

The '346 patent issued on December 20, 2005 and is entitled “Apparatus for Redundant Interconnection Between Multiple Hosts and RAID.” (JA26). The '346 patent is generally in the field of computers and systems for storing data for use by a set of computers. (JA33 at 1:7-14). In particular, the '346 patent is directed to a type of data storage system called a “RAID,” which can be accessed, for example, by a number of computers on a network (which are sometimes referred to as “host computers”). (JA35 at 5:7-26).

According to the specification, the phrase “redundant arrays of inexpensive disks (hereinafter, referred to as ‘RAID’)” defines the acronym “RAID.” (JA33 at 1:9-10; *see* JA26 at Abstract).² The '346 specification explains that “[a] RAID is a storage system based on a large capacity and a high performance, by using much

² The Abstract defines “RAID” as “redundant array of inexpensive disks,” (JA33 at 1:7-10). In Appellants’ view, the difference between array and arrays is not determinative on appeal (at least in part because an “array ... of disks” is inherently a plural construct which is not materially different from “[plural] arrays ... of disks” for purposes of comparing to what Hathorn discloses or for purposes of claim construction). The majority of extrinsic evidence defines “RAID” using a singular array, (*e.g.*, JA656, JA1434, JA1776, JA1779), as did the Board (JA189). Therefore, Appellants propose construing “RAID” as “redundant array of inexpensive disks,” as in the proceeding below (JA47).

quantity of disks, and is a fault tolerant system in which the disks or controllers etc. have a *redundant nature*.” (JA33 at 1:18-21 (emphasis added)).

The '346 patent is directed more particularly to an “apparatus for a redundant interconnection between multiple hosts and a redundant array of inexpensive disks (hereinafter, referred to as ‘RAID’), which is capable of supporting a fault tolerance of RAID controllers and simultaneously heightening a performance....” (JA26 at Abstract; JA33 at 1:7-14). The '346 patent specification distinguishes the claimed system from prior art based on the ability of the claimed system to ensure fault tolerance (meaning that a failure of a RAID controller circuit does not result in losing data (*e.g.*, JA33 at 1:35-38)) without suffering loss of performance in terms of speed.


1. The '346 Patent Specification Describes Several Prior Art Embodiments Utilizing RAID.

As Electronics and Telecommunications Research Institute (“ETRI” or “Patent Owner”) admits, “[t]he inventors of the '346 patent did not invent the concept of RAID; nor were they the first to connect multiple host computers to a RAID.” (JA133). The '346 patent illustrates several prior art systems containing RAIDs, including conventional data connections between the host computers and the RAID. (JA27-29 (FIGS. 1-3); JA33 at 1:21-2:7). Each of these prior art systems involves a storage system with high performance “in which the disks or controllers etc. have a redundant nature.” (JA33 at 1:20-21).

For example, the prior art system of FIG. 1 of the '346 patent (reproduced and annotated below) includes a RAID which has redundant controllers resulting in “twice the bandwidth and twice the performance.” (JA33 at 1:34-35).


FIG. 1

--PRIOR ART--



In this system, two host computers 100 and 101 (green) communicate with the RAID through direct connections 120 and 121 using two RAID controllers 140 and 141 (yellow), which directly control data storage operations such as reading data from and writing data to the RAID. (JA9-10; JA1-69; JA71). In the prior art system of FIG. 1, the two RAID controllers are each circuitry which enables the two host computers to transfer data to and from the disks of each RAID independently and in parallel. (JA33 at 1:24-34).


The '346 specification discloses another prior art system in FIG. 2 (reproduced and annotated below), which includes two RAID controllers 230 and 231 (yellow) which are connected to each other and with two host computers 200, 201 (green) through a “hub or switch 210” (blue) to provide for fault tolerance when one of the RAID controllers fail. (JA33 at 1:39-53).



The arrangement of this prior art configuration accounts for mistakes by the RAID controllers, and this provides so-called “fault tolerance” (*e.g.*, error recovery). (JA33 at 1:39-45). In the FIG. 2 system, each of the RAID controllers 230 and 231 has a communication controller 221 or 222, which exchanges information in the event that one or the other of the RAID controllers fail. (JA33 at 1:53-58). If one of the RAID controllers 230, 231 (yellow) fails, the hub or switch 210 (blue) will

ensure that “all of the host computers 200, 201 [(green)] are connected to a RAID controller [(yellow)] that does not have a [sic] trouble.” (JA33 at 1:50-52). The ’346 patent notes, “[h]owever, [that] in this case only a half of performance for the bandwidth ... can be obtained.” (JA33 at 1:46-59).

FIG. 3 of the ’346 patent (reproduced below with annotations) illustrates another prior art system, “which is extracted from contents disclosed in the U.S. Pat. No. 5,812,754” to Lui et al. (“Lui”). (JA33 at 1:60-2:7).




The Lui prior art system also has two RAID controllers (yellow) that are connected to host computers (green) through ports 310, 311, 320, and 321 (blue). (*Id.*). If one of the RAID controllers fails, both host computers still maintain connectivity

with the RAID. But, as in the prior art of FIG. 2, a consequence of such a failure is a reduction in performance, such as bandwidth, or in the speed of communications between the computers and the RAID. (JA33 at 2:1-7 (“this construction has no any [sic] difference from that of FIG. 2, in the structure of a communication network”)).

2. The '346 Patent Specification Describes Providing Fault Tolerance of RAID Controllers Without Sacrificing Performance.

According to the '346 patent specification, the invention provides advantages over these three prior art systems because it maintains connections without reducing performance, such as bandwidth, when a RAID controller fails. (JA33 at 2:11-15). FIG. 4 illustrates a preferred embodiment of the '346 patent (below, with annotation). This embodiment connects several host computers 400-405 (green) to two RAID controllers 460, 461 (yellow) through *two* hubs or switches 440, 441 (blue). (JA34 at 3:1-18). The “RAID” of the preferred embodiment is illustrated by a simple box 490 around the two RAID controllers 460, 461 (yellow), and the two hubs or switches 440, 441 (blue). Each RAID controller (yellow) has two Network Interface Controllers (“NIC”) (in FIG. 4, elements 470 and 471 in RAID controller 460 on the left; and elements 480 and 481 in RAID controller 461 on the right). These NICs connect each of the RAID controllers to each hub or switch (in blue).



This structure provides a “communication passage between two RAID controllers,” (JA34 at 3:64-65), that allows, *e.g.*, RAID controller 460 (yellow, on the left) to send information to RAID controller 461 (yellow, on the right) through hub/switch 441 (blue, on the right), and NIC 471 (yellow, on the left) and NIC 481 (yellow, on the right). (JA34 at 3:66-4:12). Similarly, information can be sent from RAID controller 461 (yellow, on the right) to RAID controller 460 (yellow, on the left) using hub/switch 440 (blue, on the left) and NIC 470 and 480.

In this configuration, even if one of the RAID controllers 460 and 461 (yellow) incurs trouble, the hubs or switches 440 and 441 (blue) operate such that the bandwidth remains “twice the single connection bandwidth.” (*Id.* at 3:1-9).

The specification explains that a RAID controller “having [an] error occurrence is removed from the network,” at which point a NIC from the other RAID controller “takes over a function” of a NIC on the RAID controller with the error. (JA35 at 4:19-25).

3. The '346 Patent Claims at Issue.

Independent claim 1 of the '346 patent provides:

1. An apparatus for a redundant interconnection between multiple hosts and *a RAID*, comprising:

a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers, the first RAID controlling unit including a first network controlling unit and a second network controlling unit, and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit; and

a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers, wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the numerous host computers through the plurality of connecting units, and the first network controlling unit exchanges information with the fourth network controlling unit, and the second network controlling unit exchanges information with the third network controlling unit.

(JA35 at 5:7-26 (emphases added)). Other than the mention of “a RAID” in the preamble, claim 1 does not reference or provide any additional structural or functional limitations on the claimed “RAID.” (JA35 at 5:7-8).³

Dependent claims 2-8 of the ’346 patent do not elaborate on or provide any structural restrictions on the claimed “RAID”. (JA35 at 5:27-6:19).

B. Proceedings Below

1. Petition for *Inter Partes* Review

This appeal stems from a petition filed by Petitioners on September 27, 2013, seeking IPR of the ’346 patent (docketed as IPR2013-00635). (JA36). Petitioners proposed broadest reasonable constructions, consistent with the specification, for several claim terms in the ’346 patent. (JA44). Petitioners proposed the term “RAID” be construed as “at least a redundant array of independent disks,” explaining that “[a] RAID may include RAID controllers.” (JA47 (citing JA1071-73)).⁴ The petition sought review of all claims of the ’346 patent as either anticipated or rendered obvious in light of various prior art references. (JA47-48). In particular, Petitioners asserted that U.S. Patent No.

³ Independent claim 9, which is not at issue in this appeal, similarly recites a “RAID” in the preamble without claiming any further structural or functional limitations for that term. (JA35 at 6:20-23).

⁴ Appellants have provided the joint appendix cites in place of any cites to exhibits or documents on the Board’s docket.

5,574,950 to Hathorn et al. (“Hathorn”) anticipates claims 1-3 and 5-8 of the ’346 patent based on, *inter alia*, the system depicted in Hathorn’s FIG. 3, discussed above. (JA84-99).

On January 2, 2014, Patent Owner ETRI filed its preliminary response, challenging each of Petitioners’ claim constructions and grounds for invalidating the ’346 patent claims. With respect to “RAID,” ETRI conceded that “[t]he well-known acronym itself stands for ‘redundant array of inexpensive disks,’ as noted, for example, in the ’346 Patent’s Abstract.” (JA145). ETRI criticized Petitioners’ use of the phrase “at least” in Petitioners’ construction of “RAID,” arguing that allowing “large collections of computer equipment” to constitute a RAID would be too broad. (JA145-46). ETRI proposed that “RAID” be construed as “a redundant array of independent disks, including RAID controllers, configured to be used as a peripheral by the host computers.” (JA147).

With respect to the grounds that Petitioners asserted, ETRI argued that the disclosure of, *inter alia*, Hathorn does not meet the claimed “RAID,” “RAID controllers,” and “exchanges information” limitations (as construed by ETRI). (JA151-59; JA166-73).

2. The Board’s Institution Based on Hathorn and the Institution Decision’s Construction of “RAID”.

On March 20, 2014, the Board instituted IPR proceedings on claims 1-3 and 5-8 of the ’346 patent based on anticipation by Hathorn. (JA194-99; JA204). The

Board found the broadest reasonable construction of “RAID” to be “redundant array of inexpensive disks.” (JA189). The Board explained that “‘RAID’ is well understood by a person of ordinary skill in the art as an acronym for ‘redundant array of inexpensive disks,’ and “[w]e are not persuaded that any additional restrictions are needed to capture the ordinary meaning of RAID.” (JA189 (citing JA26, Abstract)). The Board rejected both parties’ proposals as improperly and unnecessarily injecting extraneous language. (JA189).

Based on its claim constructions, the Board found that “Petitioners have made a sufficient showing of a reasonable likelihood that they would prevail with respect to their contention that Hathorn anticipates claims 1-3 and 5-8.” (JA199). The Board rejected ETRI’s argument that Hathorn’s storage controllers are not RAID controllers. (JA195-96). Instead, the Board agreed with Petitioners that “Hathorn describes a RAID configuration that can be used in connection with DASD [Direct Access Storage Device].” (JA194). Specifically, the Board concluded that there was sufficient evidence that Hathorn discloses RAID controllers because “the use of RAID is disclosed specifically in Hathorn as a type of DASD, and Petitioners’ declarant states that a person of ordinary skill in the art recognizes that a ‘DASD can be implemented in a RAID configuration.’” (JA196 (citing JA1184)). The Board also found that Hathorn disclosed each of the other

elements of claim 1, and met the additional limitations of claims 2-3 and 5-8. (JA196-98).

3. Disputes Over Construction of “RAID” Following Institution.

In its Patent Owner Response to the IPR, ETRI modified its proposed construction of, *inter alia*, “RAID.” (JA261-67). ETRI again did not dispute that “RAID” is an acronym standing for “redundant array of inexpensive disks.” (JA261). Instead, ETRI asserted that the terms “redundant,” “array,” and “disks” required further definition. ETRI argued that “disks” “can have multiple meanings” but must mean “disk drives, rather than disk platters” in the context of a RAID. (JA261). ETRI also asserted that “array” must be “a single logical storage unit of disk drives” based, not on any definitions of “array,” but rather on several proffered dictionary definitions for “RAID”. (JA261-62). ETRI argued that Hathorn supported limiting the term “RAID” to “a single logical storage unit of disk drives” because, according to ETRI, Hathorn distinguishes “RAID” and a mirroring or dual-copy system employing two disk drives. (JA262). Finally, ETRI argued that in order for a “RAID” to be “redundant” it must have a RAID controller. (JA262-63).

Based on its proposed construction of “RAID,” *inter alia*, ETRI asserted Hathorn does not anticipate because Hathorn does not disclose the claimed RAID or RAID controllers. (JA273-30).⁵

In reply, Petitioners explained that ETRI’s construction of “RAID” erroneously narrowed the broadest reasonable interpretation of the term. (JA337-40). Petitioners pointed out that neither the claim language nor the specification of the ’346 patent supports ETRI’s narrow interpretation, which imported the ambiguous notion of a “single logical unit” into the well-known term. (JA338-39). Petitioners further argued that ETRI’s narrowing of these terms conflicted with other extrinsic evidence, which define “RAID” as being “configured in many ways, either as a single unit or in various combinations of striped and mirrored configurations.” (JA339-40). Petitioners submitted the expert declaration of Dr. M. Ray Mercer in support of upholding the Board’s earlier construction of “RAID” in the Institution Decision. (JA339-40 (citing JA759; JA656; JA1072-73; JA1502-03)). Petitioners explained how Hathorn discloses a RAID as properly construed and how Hathorn’s storage controllers are the claimed “RAID controlling units.” (JA333-37).

⁵ ETRI also argued that Hathorn fails to disclose certain limitations of dependent claims 2, 3, and 5-8. (JA301-08). As discussed below, the Board did not rely on these limitations to deny Petitioner’s petition as to Hathorn.

On February 4, 2015, the Board heard oral argument from the parties. (JA368). Petitioners argued that the Board should maintain the Institution Decision's construction of RAID and reject ETRI's construction, which "packs into a RAID a number of unclaimed and extraneous limitations that make it far more narrow than the broadest reasonable interpretation." (JA374:12-15). Petitioners noted that prior art, such as Weygant, does not limit the definition of "RAID" to a single logical unit; rather, the definition is broader, "encompassing a number of different functionalities" such as the RAID Level 1 mirroring or shadowing configuration, which is what Hathorn describes. (JA374:16-375:2; JA375:8-23; JA376:14-18; JA380:1-7). Petitioners further noted that Dr. Conte, ETRI's expert, agreed that Hathorn's FIG. 3 provides mirroring. (JA382:13-18).

In response, ETRI argued that Hathorn distinguished RAID from remote copying or mirroring and that not all mirroring constitutes RAID level 1, explaining in a circular fashion that "when the disk drives being mirrored are part of a RAID array, then you have RAID Level 1 mirroring." (JA398:3-5; JA380:1-7; JA394:21-395:6). ETRI further narrowed its construction to distinguish Hathorn. ETRI asserted that RAID "must necessarily appear to the host computers as a single logical unit," which, according to ETRI, would exclude Hathorn's mirroring configuration because the DASDs in Hathorn can optionally, in one embodiment, be individually accessed by host computers. (JA400:9-15). ETRI

admitted that the intrinsic evidence itself does not restrict “RAID” to a “single logical unit” or particular RAID configuration. (JA403:21-404:21). ETRI, rather, pointed to certain extrinsic evidence in support of its construction: the 2002 Microsoft Dictionary, the 2001 Webster’s Computer Dictionary, Weygant, Chen, and statements by its declarant, Dr. Conte. (*Id.*; JA2206).

4. The Board’s Final Written Decision Narrows the Construction of “RAID” and Reverses the Initial Decision That Hathorn Likely Anticipates.

In the Final Written Decision, the Board reversed course relative to its Institution Decision in two respects: the Board changed its construction of “RAID” and, consequently, changed its anticipation analysis with respect to Hathorn. (JA5-9; JA14-23).

a. The Board’s Final Written Decision Adopted a New, Narrow Construction of “RAID” Based on Select, Partial Dictionary Definitions.

With respect to the construction of “RAID,” the Board reiterated that “‘RAID’ is well understood by a person of ordinary skill in the art as an acronym for ‘redundant array of inexpensive disks.’” (JA5). Reviewing the intrinsic evidence, the Board noted that “[t]he written description of the ’346 patent restates the acronym for RAID, but *otherwise lacks additional description of RAID or its functionality.*” (JA7 (emphasis added)). The Board *sua sponte* reasoned that “[c]onsistently and throughout the written description, RAID is referred to in the

singular” because the specification refers to “*a* RAID,” (as does claim 1) and the preferred embodiment (Figure 4) shows RAID as “a single component within a box.” (JA7-8).

The Board then turned to extrinsic evidence. The Board noted that Weygant “discloses that a RAID is a single logical unit, but also in ‘various combinations of striped and mirrored configurations,’” and that “Chen defines RAID to be Redundant Arrays of Inexpensive Disks.” (JA8 (emphasis and underline in original)). The Board reasoned that the specification, Weygant, Chen, and Dr. Conte’s testimony “support that a RAID is a single logical unit.” (JA8). To support this new construction, the Board relied on two dictionary definitions that admittedly were “published after the foreign priority date of the ’346 patent...[to] substantiate the proposition that those of ordinary skill generally refer to RAID as a single logical unit.” (JA8-9).

With respect to Weygant’s broader disclosure of a RAID, the Board explained that “[a]lthough Weygant also indicates that a RAID can exist in forms other than a singular logical unit, those forms are limited to ‘striped and mirrored’ configurations.” (JA8). The Board expressly excluded the striped and mirrored forms of a RAID from the scope of this new construction because “[w]e have not been shown evidence that the RAID of the ’346 patent is configured in a ‘striped and mirrored configuration,’ which might not be in the single unit configuration of

RAID.” (JA9). Based on this reasoning, the Board modified its original construction of RAID to “a single logical unit for mass storage using multiple physical disk drives.” (JA9).

b. The Board’s Final Written Decision Relied on This New Narrowed Definition of “RAID” to Conclude that Hathorn Does Not Anticipate the ‘346 Claims.

The Board applied this new, narrowed construction of “RAID” to find that Hathorn does not satisfy the “RAID” requirement of the ’346 patent claims at issue. In response to Petitioner’s arguments that the Hathorn disclosure teaches that Hathorn’s “DASDs [Direct Access Storage Devices] can be *configured* as a RAID,” the Board stated that “[b]eing capable of a RAID configuration is not the same as an actual RAID” and that evidence supported distinguishing a DASD from the construed RAID. (JA22). The Board also rejected Petitioner’s argument “‘that DASDs can be arranged as a RAID.’” (*Id.*).

The Board admitted that “the record does support that RAID level 1, one of several possible RAID configurations, is ‘disk mirroring,’” (*Id.*). The Board further admitted that “*Hathorn may disclose a RAID level 1 configuration....*” (*Id.* (emphasis added)). However, the Board framed “the question before us [a]s whether Hathorn discloses RAID as we have construed the term, i.e., a ‘single logical unit’.” (JA22-23). The Board answered its own question in the negative.

Using its new, narrowed construction of RAID, the Board found that Hathorn's direct access storage devices "are not a single logical unit because primary host 301 *can* individually (or directly) access either one of those." (JA23 (emphasis in original)). To support this conclusion, the Board did not cite evidence in the record, but rather exclusively cited an argument made by ETRI's counsel at oral hearing. (JA23 (citing JA417:4-7)). For those reasons, the Board found that Petitioners had not shown by a preponderance of evidence that Hathorn discloses a RAID as the Board construed that term. (*Id.*).

The Board did not address further arguments about patentability, and ordered that "claims 1-3 and 5-8 of U.S. Patent No. 6,978,346 have not been shown by a preponderance of the evidence to be unpatentable." (JA24).

C. Hathorn

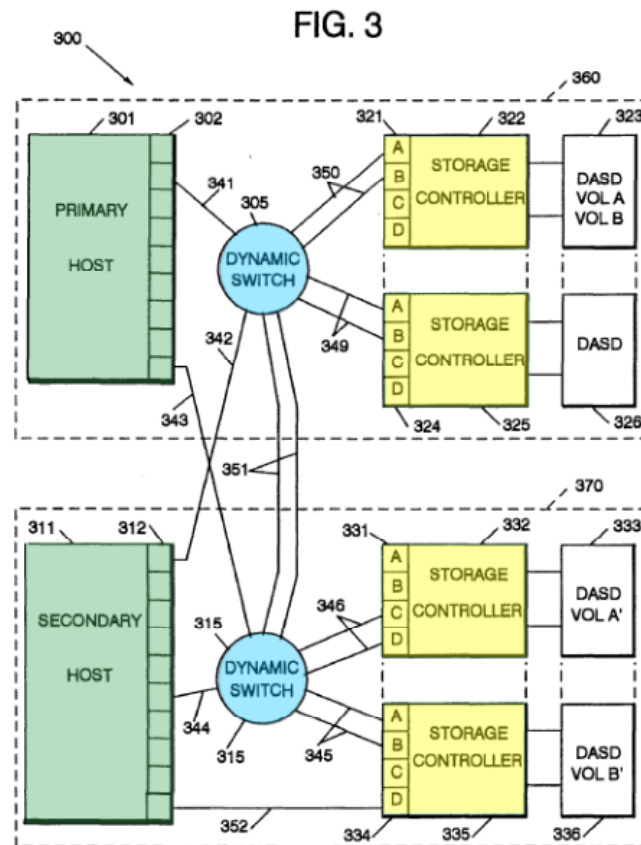
As discussed in § I.B.2, *supra*, the Board instituted this IPR of claims 1-3 and 5-8 based on Hathorn. (JA1033-50).

Hathorn, titled "Remote Data Shadowing Using Multimode Interface to Dynamically Reconfigure Control Link-Level and Communication Link-Level," issued on November 12, 1996, and lists IBM as the assignee. (JA1033). Hathorn discloses an improved system for connecting and communicating between computers and storage controllers. (JA1044 at 4:39-45). Hathorn explains that "data storage usually involves direct access storage devices (DASD). DASD

storage, for example, can comprise magnetic and/or optical disks.” (JA1043 at 1:28-31). “Having a back-up data copy is mandatory for many businesses as data loss could be catastrophic to the business.” (JA1043 at 1:60-61). Such back-up data can include dual copy. “An example of dual copy involves providing additional DASD’s so that data is written to the additional DASDs (sometimes referred to as *mirroring*).” (JA1043 at 1:63-67 (emphasis added)). Hathorn provides that “[a]nother data back-up alternative that overcomes the need to double the storage devices involves writing data to a *redundant array of inexpensive devices (RAID) configuration*” and that “[c]urrently *there are several different RAID configurations available*.” (JA1043 at 2:4-11 (emphases added)).

Hathorn further teaches that “data disaster recovery solutions include remote dual copy wherein data is backed-up not only remotely, but also continuously.” (JA1044 at 4:1-3). According to Hathorn, “[i]n a typical remote dual copy system, there may exist multiple primary processors connected ... to multiple primary storage controllers, each having strings of primary DASDs attached thereto” and that “[a] similar processing system may exist at a remote secondary site. Additionally, many communication links may be required to connect primary processors to secondary processors and/or secondary storage controllers, and primary storage controllers may be connected to secondary storage controllers and/or secondary processors.” (JA1044 at 4:12-21).

Hathorn discloses a “remote copy” system with dynamically modifiable ports on the storage controller. (JA1033 at Abstract). Using these modifiable ports, “shared communication links . . . can dynamically interface either a host processor to a storage controller, or can interface one storage controller to another storage controller.” (JA1044 at 4:30-37). FIG. 3 depicts a preferred embodiment of Hathorn’s system (reproduced below with highlights added).



In FIG. 3, there are two “host” computers: a primary host computer 301 and a secondary host computer 311 (highlighted in green). (JA1036 (FIG. 3)). The system also contains primary storage controllers 322 and 325 and secondary

storage controllers 332 and 335 (highlighted in yellow). (*Id.*). These storage controllers control the access to and from the disks of the Direct Access Storage Disks (DASD) which store data needed by the computer network, including “data or records for back-up” to the host computers. (JA1046 at 7:28-44). Two dynamic switches 305 and 315 (highlighted in blue) connect the primary host computer, secondary host computer, primary storage controllers and secondary storage controllers. These dynamic switches enable communications between these components of the system. (JA1046 at 7:14-26, 7:59-62, 8:3-14). Hathorn explains that this system performs “actual data shadowing” (in other words, backup copying) which “involves the primary storage controller 325 sending data to be shadowed to the secondary storage controller 335.” (JA1047 at 9:29-51).

D. Extrinsic Evidence from the Record Below

The following provides a brief description of each piece of extrinsic evidence which the parties and/or Board relied on to establish the plain and ordinary meaning of “RAID” to a person of ordinary skill at the time of the invention in 2000:

1. Prior Art References

a. Weygant

Peter Weygant authored *Clusters for High Availability: A Primer of HP-UX Solutions*, a Hewlett-Packard Professional Book from 1996 (hereinafter

“Weygant”). (JA592-596). In describing an architecture for a cluster of computers using the HP-UX operating system, Weygant discusses methods for eliminating disks as single points of failure by, *e.g.*, using disk arrays in RAID configurations. Weygant explains the basic terminology of RAIDs and the different “levels” or conceptual configurations for a RAID:

The acronym RAID stands for *redundant array of inexpensive disks*. A group of disks function together in a variety of configurable arrangements known as RAID levels. Some levels allow hardware mirroring, while others provide protection through the use of parity data, which allows the array to reconstruct lost data if a disk mechanism fails.

Common RAID levels are as follows:

- Level 0: the controller writes data to all disks in stripes.⁶ This level provides no data protection.
- Level 1: the controller writes data to *mirrored groups of disks*.
- Level 3: data is striped byte-wise, and the controller stores parity information on a separate disk so that lost data from any disk can be recovered.
- Level 5: data is striped block-wise, and the controller spreads parity information across all disks so that lost data from any disk can be recovered.

(JA656 (emphases added)). In a “Glossary of High Availability Terminology,”

Weygant similarly states:

⁶ Data “stripes” are segments of data which are logically sequential, but which are then physically broken up and stored on different physical storage devices, for example across multiple disks in a RAID. This “striping” allows logically connected data to be retrieved more quickly because the segments can be retrieved in parallel, as opposed to sequentially if they were all stored on just one disk or device. (JA1445).

RAID is an acronym for *redundant array of inexpensive disks*. A RAID device consists of a group of disks that can be configured in many ways, *either as a single unit or in various combinations of striped and mirrored configurations*. The types of configuration available are called RAID levels:

- RAID 0: Disk striping
- RAID 1: *Disk mirroring*.
- RAID 0/1: Sector Interleaved groups of mirrored disks. Also called RAID 1/0 or RAID 10
- RAID 2: Multiple check disks using Hamming code.
- RAID 3: Byte striped, single check disk using parity.
- RAID 4: Block striped, single check disk using parity.
- RAID 5: Block striped, data and parity spread over all disks.

(JA759 (emphases added)).

b. Chen

“RAID: High-Performance, Reliable Secondary Storage,” authored by Chen et al. (hereinafter, “Chen”) is an article that focuses on RAID. (JA1435).⁷ Chen provides “a comprehensive overview of disk arrays and provides a framework in which to organize current and future work.” (JA1435).

Chen explains that interest in RAID technology had been increasing due to exponential improvements in semiconductor technology that made microprocessors faster, and enabled larger memory systems “which in turn require larger, higher-performance secondary storage systems.” (JA1438). Chen posits that “[d]isk arrays, which organize multiple independent disks into a large, high-

⁷ Patent Owner’s expert admitted that Chen was considered an authoritative document in this field at the time of the invention. (JA2105:25-JA2106:17).

performance logical disk, are a natural solution to the problem” but are “highly vulnerable to disk failures.” (JA1439). In response to this vulnerability, “different data striping and redundancy schemes have been developed,” including seven known RAID configurations. (JA1439). Chen “describes the basic RAID, Redundant Arrays of Inexpensive Disks, organizations [sic] which will be used as the basis for further examinations of performance, cost and reliability of disk arrays.” (JA1446). Chen also depicts how disks are configured for each RAID level:



(JA1448).⁸ In each illustrated RAID Level, Chen shows a user capacity of four disks, with the shaded portions representing the redundant or copied information for backup storage purposes. (JA1448). Chen provides detailed explanations of

⁸ Chen explains that “[s]trictly speaking RAID Level 0 is not a type of *redundant* array of inexpensive disks since it stores no error-correcting codes.” (JA1447, n.1).

each of these RAID levels. For example, with respect to “Mirrored (RAID Level 1),” Chen states:

The traditional solution, called *mirroring* or *shadowing*, uses twice as many disks as a nonredundant disk array [Bitton88]. Whenever data is written to a disk the same data is also written to a redundant disk, so that there are always two copies of the information. When data is read, it can be retrieved from the disk with the shorter queueing, seek and rotational delays [Chen90a]. If a disk fails, the second copy is used to service requests. Mirroring is frequently used in database applications where availability and transaction rate are more important than storage efficiency [Gray90].

(JA1447 (emphasis in original)).

2. Dictionaries

a. Microsoft Computer Dictionary (1999)

The fourth edition of the Microsoft Computer Dictionary was published in 1999. (JA1433). It provides the following definition of the term “RAID”:

RAID \rad\ n. *Acronym for redundant array of independent disks (formerly called redundant array of inexpensive disks). A data storage method in which data, along with information used for error correction, such as parity bits or Hamming codes, is distributed among two or more hard disks in order to improve performance and reliability. The hard disk array is governed by array management software and a disk controller, which handles the error correction. RAID is generally used on network servers. Several defined levels of RAID offer differing trade-offs among access speed, reliability, and cost. See also disk controller, error-correction coding, Hamming code, hard disk, parity bit, server (definition 1).*

(JA1434 (emphasis added)).

b. Webster's Computer Dictionary (2001)

The ninth edition of Webster's New World Computer Dictionary was published in 2001. (JA1775). That dictionary defines "RAID" as follows:

Acronym for Redundant Array of Inexpensive Disks or Redundant Array of Independent Disks. A group of hard disks under the control of array-management software that work together to improve performance and decrease the odds of losing data due to mechanical or electronic failure by using such techniques as data striping. Because of their complexity and cost, RAID implementations are most often used on network servers. Several RAID levels exist, each with advantages and disadvantages. RAID arrays are generally used for high-volume servers. See RAID 0, RAID 1, RAID 2, RAID 5, RAID 10.

(JA1776 (emphasis added)). Webster's then details various configurations of RAID. (JA1776). For example, Webster's defines "RAID 1" as follows:

A type of RAID storage device that combines two or more hard disks into a single logical drive, but—in contrast to RAID 0—in a way that backs up the data so that nothing is lost if one of the drives should fail. Performance is sacrificed for the sake of data integrity. See RAID.

(JA1776). The definitions for RAID 0, 2, and 5 similarly begin by describing "[a] type of RAID storage device that combines two or more hard disks into a single logical drive...." (JA1776).

c. Microsoft Computer Dictionary (2002)

In 2002, Microsoft published the next edition of the Microsoft Computer Dictionary. (JA1778). This 2002 dictionary contains slight variations to the definition of RAID (compared to the fourth edition discussed above):

Acronym for redundant array of independent (or inexpensive) disks. A data storage method in which data is distributed across a group of computer disk drives that function as a single storage unit. All the information stored on each of the disks is duplicated on other disks in the array. This redundancy ensures that no information will be lost if one of the disks fails. RAID is generally used on network servers where data accessibility is critical and fault tolerance is required. There are various defined levels of RAID, each offering differing trade-offs among access speed, reliability, and cost. See also disk controller, error-correction coding, Hamming code, hard disk, parity bit, server (definition 1).

(JA1779 (emphases added)).

II. SUMMARY OF ARGUMENT

This appeal stems from two errors in the Board’s Final Written Decision.

First, the Board improperly construed “RAID” by selectively using extrinsic evidence that was contrary to the intrinsic evidence and the plain meaning of the term to a person of ordinary skill at the time of the alleged invention. The proper construction of “RAID” should be, as the Board found in the Institution Decision, a “redundant array of inexpensive disks” (JA189), and *not* “a single logical unit for mass storage using multiple physical disk drives,” as the Board held in its Final Written Decision (JA8-9).

ETRI admits that the ’346 inventors did not invent the concept of a RAID. (JA133 at 2). This was well-known technology described by a well-known acronym by the time of the filing of the application that led to the ’346 Patent. The intrinsic evidence confirms that the inventors intended the term to have its well-

known and straightforward meaning. The '346 patent repeatedly describes “RAID” as a “redundant array of inexpensive disks” and explains that a “RAID is a storage system based on a large capacity and a high performance, by using much quantity of *disks* [sic], and is a fault tolerant system in which the disks or controllers etc. *have a redundant nature*.” (JA26 at Abstract; JA33 at 1:7-10; JA33 at 1:18-21 (emphases added)). Similarly, each and every piece of extrinsic evidence in the record defines “RAID” as “redundant array of inexpensive [or independent⁹] disks” and describes several RAID configurations encompassed by the term. (JA1434; JA759; JA1446; JA1776; JA1779). Thus the broadest reasonable interpretation, and the plain and ordinary meaning of “RAID” to a person of ordinary skill is “redundant array of inexpensive disks.”

Largely guided by ETRI’s attorney argument, the Board ignored the intrinsic evidence and cherry-picked through the extrinsic evidence to narrow the term beyond its plain and ordinary meaning. The Board imported the phrase “single logical unit” into its construction—a phrase that does not appear anywhere in the intrinsic *or* extrinsic evidence—all the while eliminating mention of “redundant,” which the specification uses in every description of RAID and the extrinsic evidence echoes. (JA8). In support, the Board made legally erroneous conclusions

⁹ References refer to the “I” in “RAID” to mean either “inexpensive” or “independent,” but this term does not affect the issues before the Court in this appeal.

about the intrinsic evidence, reasoning that “a RAID” necessarily must be singular” (JA7-8) and requiring the specification to describe configurations known to be encompassed by the term “RAID” (JA8-9). And, the Board misread the extrinsic evidence—*e.g.*, citing Weygant for support for “single logical unit” when it actually uses the term “single unit” and defines “RAID” more broadly to include configurations other than “a single unit”—while ignoring that each definition in the extrinsic evidence defines the term “RAID” more broadly than the Board ultimately did. (JA8-9).

Moreover, the Board failed to appreciate that even if the extrinsic evidence supported a narrower construction, that alone cannot provide the basis for narrowing a claim term. *3M Innovative Prop. Co. v. Tredegar Corp.*, 725 F.3d 1315, 1333-34 (Fed. Cir. 2013); *Tegal Corp. v. Tokyo Electron Am., Inc.*, 257 F.3d 1331, 1345 (Fed. Cir. 2001). The Board further narrowed “single logical unit” to exclude RAID level 1 configurations (which all the extrinsic evidence included) and to exclude any configurations where the host computer accesses storage controllers individually or directly (supported only by ETRI’s attorney argument in efforts to avoid anticipation). These errors, separately and collectively, warrant reversal of the Board’s claim construction of “RAID” and a return to its initial construction of “RAID” as “redundant array of inexpensive disks.”

Second, these claim construction errors bled into the Board’s anticipation analysis. The Board relied on the requirements that “RAID” be a “single logical unit” that excludes RAID Level 1 mirroring and systems where the host directly accesses storage controllers to reject Hathorn as anticipating claims 1-3 and 5-8 of the ’346 patent. (JA22-23). But even this analysis was flawed; Hathorn satisfies “RAID” as ultimately construed by the Board because Hathorn discloses alternative embodiments in which direct access does not occur. And under the proper interpretation of RAID, there is no dispute that Hathorn meets all the claim limitations, including “RAID”—the Board, in its Final Written Decision, even admitted that “Hathorn may disclose a RAID level 1 configuration.” (JA22-23).

Thus, Appellants respectfully request that the Court reverse the Board’s claim construction of “RAID” and its finding that Hathorn does not anticipate claims 1-3 and 5-8 of the ’346 patent.

III. ARGUMENT

A. Standard of Review.

The Board must give claims their broadest reasonable construction consistent with the specification. 37 C.F.R. § 42.100(b); *In re Cuozzo Speed Techs., LLC*, 793 F.3d 1268, 1277-79 (Fed. Cir. 2015). This Court reviews the Board’s claim construction according to the Supreme Court’s decision in *Teva Pharmaceuticals USA, Inc. v. Sandoz, Inc.*, 135 S. Ct. 831, 841 (2015). *See In re*

Cuozzo Speed Techs., LLC, 793 F.3d 1268, 1280-1281 (Fed. Cir. 2015) (applying *Teva* in an appeal from *inter partes* review). Under *Teva*, underlying factual determinations concerning extrinsic evidence are reviewed for substantial evidence, while the Board’s determinations based on the intrinsic evidence and the Board’s ultimate construction of the claim are reviewed *de novo*. *Teva*, 135 S. Ct. at 841.

Anticipation is a question of fact, reviewed for substantial evidence on appeal. *In re Baxter Travenol Labs.*, 952 F.2d 388, 390 (Fed. Cir. 1991). “Substantial evidence is something less than the weight of the evidence but more than a mere scintilla of evidence ... and means such relevant evidence as a reasonable mind might accept as adequate to support a conclusion.” *In re Suitco Surface, Inc.*, 603 F.3d 1255, 1260 (Fed. Cir. 2010) (internal citations and quotations omitted). This Court has vacated anticipation decisions when the Board’s “broadest reasonable construction” was unreasonable or inconsistent with the patent’s claim language or the specification. *Id.* at 1259-61.

B. The Board’s Narrowed Construction of “RAID” is Inconsistent with the ’346 Patent and How a Person of Ordinary Skill in the Art Would Understand the Term.

The Board and the parties agree that the term “RAID” does not carry any special meaning within the context of the ’346 patent, and thus must be given its “ordinary meaning as understood by persons skilled in the art in question at the

time of the invention.” *InterDigital Comm’ns, LLC v. ITC*, 690 F.3d 1318, 1324 (Fed. Cir. 2012). Neither the written description nor the prosecution history of the patent indicates that the patentee acted as his own lexicographer or that the term “RAID” has any different meaning than what is expressly stated in the specification and is understood by one of ordinary skill in the art at the time of invention. Thus, this claim construction boils down to what the plain and ordinary meaning is in the context of the BRI standard.

The Board properly construed “RAID” in the Institution Decision as a “redundant array of inexpensive disks,” which is the specification’s definition of the acronym “RAID.” (JA189; JA26 at Abstract; JA33 at 1:7-10). As the Board originally ruled, “‘RAID’ is well understood by a person of ordinary skill in the art as an acronym for ‘redundant array of inexpensive disks’.... We are not persuaded that any additional restrictions are needed to capture the ordinary meaning of RAID.” (JA189 (emphases added)). Indeed, this construction is most consistent with the specification and all of the contemporaneous prior art and dictionaries defining the term “RAID.” But in its Final Written Decision, the Board erroneously narrowed the construction of RAID based solely on its cherry-picking of certain extrinsic evidence and misreading of the specification to require “a single logical unit for mass storage using multiple physical disk drives.” (JA9).

There is virtually no support for considering that construction to be the plain and ordinary meaning of or broadest reasonable interpretation of “RAID.”

1. The Broadest Reasonable Interpretation of “RAID” Is “Redundant Array of Inexpensive Disks”

The proper construction of “RAID” is the plain and ordinary meaning of that term: “a redundant array of inexpensive disks,” as set forth in the explicit language of the specification and as is consistent with all of the extrinsic evidence. (*See* JA189). The Board’s overly narrow construction of “RAID” as “a single logical unit for mass storage using multiple physical disk drives” explicitly excludes several concepts and configurations encompassed by the proper construction set forth in the specification, while importing unsupported restrictions on “RAID”.

a. The Intrinsic Evidence Supports Construing “RAID” to Mean a “Redundant Array of Inexpensive Disks”

The intrinsic record supports “redundant array of inexpensive disks” as the broadest reasonable interpretation of “RAID.” The specification is “the single best guide to the meaning of a disputed term.” *MBO Labs., Inc. v. Becton, Dickenson & Co.*, 474 F.3d 1323, 1329 (Fed. Cir. 2007) (quoting *Phillips v. AWH Corp.*, 415 F.3d 1303, 1315 (Fed. Cir. 2005)).

Here, the ’346 specification repeatedly provides the same broad definition of “RAID”: redundant array[s] of inexpensive disks. (JA33 at 1:7-10; JA26 at Abstract; *see* JA145). The ’346 patent specification also emphasizes that

redundancy is a key attribute of RAID. The “R” in “RAID” stands for “redundant.” (JA33 at 1:7-10). The specification explains that “RAID is a storage system based on a large capacity and high performance, by using much quantity of disks [sic], and is a fault tolerant system in which the disks or controllers etc. *have a redundant nature.*” (JA33 at 1:18-21 (emphasis added)).

Nothing in the ’346 patent specification “demonstrate[s] an intent to deviate from the ordinary and accustomed meaning of [RAID] by redefining the term or by characterizing the invention in the intrinsic record using words or expressions of manifest exclusion or restriction.” *Teleflex, Inc. v. Ficosa N. Am. Corp.*, 299 F.3d 1313, 1324 (Fed. Cir. 2002). In particular, there is nothing in the intrinsic evidence to support narrowing the construction. The intrinsic record contains no limitations or functional requirements of a specific RAID configuration covered by the ’346 patent. None of the ’346 patent figures illustrates particular configurations of disks in a RAID in a way that excludes other configurations; RAID is merely depicted as a box in the various exemplary and prior art embodiments. (*E.g.*, JA27-32). And the one exemplary configuration of RAID shown in FIG. 4 is “at least” a RAID Level 1 mirrored configuration—a configuration which the Board incorrectly excludes from the claim scope (JA22-23). Thus, the intrinsic record supports a broad construction of RAID that includes well-known redundant configurations. *See InterDigital*, 690 F.3d at 1324 (refusing to narrow claim term where “[n]either

the specification nor the prosecution history contains a restrictive definition ... and the patentee did not at any point disavow the broader interpretation of the that term”).

None of this intrinsic evidence is disputed. In fact, ETRI’s expert, Dr. Conte, confirmed that a person of ordinary skill in the art reviewing the ’346 patent specification would come to the same conclusion. He agreed that the ’346 patent’s definition of “RAID” as a “redundant array of inexpensive disks” is correct. (JA2065:15-18 (“I think that definition is clearly correct since it’s an expansion of the acronym.”)). Dr. Conte admitted that the ’346 patent specification does not depict or describe the number of disks in a RAID or require a particular RAID configuration. (JA2056:3-18). He explained that it would be understood to one skilled in the art “that there are disk drives in the RAID and they’re not illustrated.” (JA2059:10-12). And, significantly, Dr. Conte testified that Figure 4, which is an exemplary embodiment, depicts “at least RAID 1,” which is known as mirroring. (JA2057:21-25).

The Board similarly noted that “‘RAID’ is well understood by a person of ordinary skill in the art as an acronym for ‘redundant array of inexpensive disks.’” (JA189). It observed that “[t]he written description of the ’346 patent restates the acronym for RAID, but otherwise lacks additional description of RAID or its functionality.” (JA7). Thus, both the Board and all the parties agree that the

intrinsic evidence supports defining RAID as “redundant array of inexpensive disks,” that the specification does not provide any further limitations on “RAID,” and that the specification does not limit RAID in the ’346 patent to a particular level or configuration. This should have ended the Board’s inquiry: the broadest reasonable construction of “RAID” is “redundant array of inexpensive disks.”

But it did not. The Board opted to further narrow the construction of “RAID” to “a single logical unit for mass storage using multiple physical disk drives.” This construction—unlike the construction adopted by the Board in its Institution Decision and advocated by Petitioners—contradicts the intrinsic record in several respects and derives entirely from the extrinsic record and ETRI’s attorney argument. The Board did so, despite this Court’s warning that “undue reliance on extrinsic evidence poses the risk that it will be used to change the meaning of claims in derogation of the ‘indisputable public records consisting of the claims, the specification, and the prosecution history,’ thereby undermining the public notice function of patents.” *Nystrom v. TREX Co.*, 424 F.3d 1136, 1143 (Fed. Cir. 2005) (quoting *Southwall Techs., Inc. v. Cardinal IG Co.*, 54 F.3d 1570, 1578 (Fed. Cir. 1995)).

First, the Board eliminates “redundant” from its construction of “RAID.” This conflicts with the specification’s explicit and repeated requirement that RAID be redundant. (JA33 at 1:18-21, 1:7-10). Indeed, even ETRI’s expert focused on

the '346 patent's redundancy requirement in discussing his understanding of RAID. (JA2057:9-58:12).

Second, the Board relies on the "single logical unit" phrase to read out RAID level 1 mirroring from the meaning of "RAID," all the while admitting that "[t]he record does support that RAID level 1, one of several possible RAID configurations, is 'disk mirroring.'" (JA22). To justify this narrowing of "RAID," the Board explains, in part, that "[w]e have not been shown evidence that the RAID of the '346 patent is configured in a 'striped and mirrored configuration,' which might not be in a single unit configuration of RAID." (JA9). This statement turns the fundamental canons of claim construction on their head. The intrinsic record can only prescribe a more particularized, narrower meaning of a term: "[i]t is axiomatic that we will not narrow a claim term beyond its plain and ordinary meaning unless there is support for the limitation in the words of the claim, the specification, or the prosecution history." *3M Innovative Props.*, 725 F.3d at 1333 (citations omitted). Patent law does not require the specification to list or define that which is already well known. *See Streck, Inc. v. Research & Diagnostic Sys., Inc.*, 665 F.3d 1269, 1285 (Fed. Cir. 2012). The intrinsic record broadly describes RAID and is devoid of any statements or limitations that exclude RAID level 1 mirroring from the scope of the '346 patent, and, excluding RAID 1 from the term "RAID" would be contrary to what a person of ordinary skill would understand

Figure 4, the '346 patent's preferred embodiment, to cover. (JA2057:21-25 (“Q...Can you tell from looking at the figure or from reading the specification what level RAID we have in Figure 4? A: Let me look at the specification. I’m able to determine it’s at least RAID 1.”)). Reading out embodiments “is rarely, if ever, correct.” *InterDigital*, 690 F.3d at 1326 (internal quotations omitted) (citations omitted).

Third, there is no support in the intrinsic record for limiting “RAID” to “a single logical unit.” This limitation is, at best, entirely imported from the extrinsic record, which is legally improper. *See 3M Innovative Props.*, 725 F.3d at 1333-34. The term “single logical unit” does not once appear in the intrinsic record and neither the Board nor ETRI identified any disclosures in the intrinsic record that describe a single logical unit. (*See* JA338; JA340; JA2066:21-2067:5). This Court has warned against importing limitations that do not appear in the intrinsic record. *See 3M Innovative Props.*, 725 F.3d at 1333 (refusing to inject specific size requirements provided in extrinsic evidence into construction of “ribbon” where no support existed in the intrinsic record) (citations omitted).

The Board’s only cited support in the specification is the use of “a” to describe “a RAID” in the specification and claims—an argument that was never presented by either party. (JA7-8). The Board reasons that “[c]onsistently and throughout the written description, RAID is referred to in the singular,...[t]he

claims also recite ‘a RAID.’” (JA7). The Board also notes that a preferred embodiment depicted in Figure 4 of the ’346 patent “shows RAID 490 as a single component within a box.” (JA8). The implication of this argument appears to be that “a RAID” requires RAID to be singular, which supports the construction of the term as “a singular logical unit.” Notwithstanding the fact that the Board’s leap from singular RAID necessarily requiring a singular logical unit strains logic, regardless, it is contrary to law. “As a general rule, the words ‘a’ or ‘an’ in a patent claim carry the meaning of ‘one or more.’” *01 Communique Lab., Inc. v. LogMeIn, Inc.*, 687 F.3d 1292, 1297 (Fed. Cir. 2012) (quoting *TiVo, Inc. v. EchoStar Commc’ns Corp.*, 516 F.3d 1290, 1303 (Fed. Cir. 2008)) (emphasis added); *Baldwin Graphic Sys., Inc. v. Siebert, Inc.*, 512 F.3d 1338, 1342 (Fed. Cir. 2008)).¹⁰ “The exceptions to this rule are extremely limited: a patentee must evidence a clear intent to limit ‘a’ or ‘an’ to ‘one’.” *Id.* Here, the Board does not even attempt to identify any such intrinsic evidence. Further, as noted above, the acronym RAID stands for an “array of ... disks,” which is inherently a construct made of plural elements. Therefore, the Board’s emphasis of “a” singular RAID ignores the plain language of the term, which is a plural construct. Indeed, neither the term nor the specification provides a basis to distinguish between one array of

¹⁰ And notably, even if Figure 4 or any other embodiments depict a single RAID, that, alone, is insufficient to limit the term without more. *See 3M Innovative Props.*, 725 F.3d at 1321.

multiple disks versus multiple arrays of multiple disks (because such line drawing, implicit in the Board’s ruling, is meaningless and arbitrary given that whether one is talking about one array or multiple arrays, the arrays are always made up of plural disks). Thus, if anything, the use of “a RAID” in the specification and claims signals that the ’346 patent more broadly covers one or more RAIDs, while lending zero support for the Board’s claim construction, particularly under the BRI standard. *See* § III.A, *supra*.

Thus, the intrinsic record provides clear support for construing “RAID” consistent with its plain and ordinary meaning, which is a “redundant array of inexpensive disks” and which includes various redundant configurations. By contrast, it lends no support—and contradicts—the Board’s construction.

b. Extrinsic Evidence Confirms That the Plain And Ordinary Meaning of “RAID” Is a “Redundant Array of Inexpensive Disks”¹¹

“A court may look to extrinsic evidence so long as the extrinsic evidence does not contradict the meaning otherwise apparent from the intrinsic record.” *Helmsderfer v. Bobrick Washroom Equip., Inc.*, 527 F.3d 1379, 1382 (Fed. Cir. 2008) (citation omitted). The definitions of “RAID” in dictionaries and prior art

¹¹ Petitioners believe the intrinsic evidence is sufficient to end the claim construction inquiry and that resort to extrinsic evidence is unnecessary. *See W.E. Hall Co. v. Atlanta Corrugating, LLC*, 370 F.3d 1343, 1350 (Fed. Cir. 2004); *see also Tempo Lighting, Inc. v. Tivoli, LLC*, 742 F.3d 973, 977 (Fed. Cir. 2014). However, because the Board and Patent Owner relied on extrinsic evidence, Petitioners discuss such evidence here to ensure the record is complete.

references contemporaneous with the '346 patent serve to further confirm that a person of ordinary skill in the art at the time of the alleged invention would understand that “RAID” broadly refers to a “redundant array of inexpensive disks,” inclusive of known RAID configurations.

Each dictionary and prior art reference at issue defines “RAID” slightly differently, but they all repeat common attributes and descriptions of “RAID” disclosed in the '346 patent specification. Each dictionary and prior art reference defines “RAID” as an acronym for “redundant array of inexpensive [or independent] disks”—just like the '346 patent. (JA1434; JA759; JA1446; JA1776; JA1779; *see also* JA1502 at ¶ 6).

Moreover, each dictionary and prior art reference explains that there are several levels or configurations of RAID. For example, both Microsoft dictionaries state that there are “[s]everal defined levels of RAID [which] offer differing trade-offs....” (JA1434; *see* JA1779). And, Webster’s Computer Dictionary identifies several possible configurations for RAID and explains that “[s]everal RAID levels exist, each with advantages and disadvantages.” (JA1776). Where there is “more than one dictionary definition [that] is consistent with the use of the words in the intrinsic record, the claim terms may be construed to encompass all such consistent meanings,” *Texas Digital Sys., Inc. v. Telegenix, Inc.*, 308 F.3d 1193, 1203 (Fed. Cir. 2002), and here, that construction would

encompass all the redundant configurations of RAID disclosed in these dictionaries.

Notably, the Board exclusively relied on Webster's Computer Dictionary and the 2002 Microsoft Computer Dictionary, without regard to other dictionary definitions; in doing so, it committed legal error by focusing exclusively on post-priority date dictionary definitions in the face of other contemporaneous evidence. *See Starhome GmbH v. AT&T Mobility LLC*, 743 F.3d 849, 856 (Fed. Cir. 2014) ("The term 'gateway' had a well-understood meaning in the art *at the time the patentees filed the application that led to the '487 patent*. As evidenced by technical dictionaries, one of ordinary skill in the art would have understood a 'gateway' to be a connection between different networks." (emphasis added) (citations omitted)); *Mass. Inst. of Tech. & Elec. For Imaging, Inc. v. Abacus Software*, 462 F.3d 1344, 1357 (Fed. Cir. 2006) (same). Nevertheless, the dictionary definitions the Board relies on do not mandate a different claim construction that would alter or conflict with the broadest reasonable construction of "RAID"—as "redundant array of inexpensive disks" that encompasses all redundant configurations.

Similarly, Weygant explains that "[a] RAID device consists of a group of disks that can be configured in many ways..." and describes various RAID levels. (JA759). Weygant recognizes the known mirroring or shadowing configurations

of a RAID by labeling “RAID 1: Disk *Mirroring*” and explains that in RAID Level 1, “the controller writes data to *mirrored* groups of disks”. (JA759 (emphasis added); JA656 (emphasis added)).

The Chen reference, which ETRI’s expert admitted was an authoritative publication at the time of the alleged invention, similarly provides a detailed description and depiction of various RAID levels. (JA1446-47; *see* JA2105:25-2106:17). In particular, Chen describes “*Mirrored* (RAID Level 1)” as “the traditional solution, called *mirroring* or *shadowing*.” (JA1447 (first emphasis added, second emphasis in original)).

Thus, all the extrinsic evidence confirms that “RAID” is understood by persons of ordinary skill in the art to mean “redundant array of inexpensive disks” and to encompass various redundant configurations. And the repeated disclosure of, *e.g.*, RAID Level 1 mirroring or shadowing as a well-known redundant configuration in these references is consistent with the ’346 specification (which discloses to skilled artisans a preferred embodiment that includes “at least” RAID Level 1 mirroring). This extrinsic evidence demonstrates that, contrary to the Board’s construction of “RAID” (JA22-23), RAID Level 1 mirroring would necessarily be encompassed within the plain and ordinary meaning of “RAID” to a person of ordinary skill in the art.

c. The Extrinsic Record Does Not Support the Board's Construction.

The Board erred first and foremost by focusing primarily on extrinsic rather than intrinsic evidence. *See Power Integrations, Inc. v. Fairchild Semiconductor Int'l, Inc.*, 711 F.3d 1348, 1362 (Fed. Cir. 2013) (“our inquiry here starts with the intrinsic record, including the specification, and not with a dictionary definition of the disputed term”); *Combined Sys., Inc. v. Defense Tech. Corp. of Am.*, 350 F.3d 1207, 1215 (Fed. Cir. 2003) (error to place primary importance on extrinsic evidence and allow it to overcome the clear language of the patent itself). But that was not the Board’s only error. The Board also improperly turned a blind eye to common disclosures throughout the extrinsic evidence that are consistent with the intrinsic record and support Petitioners’ construction. The Board chose to rely on certain extrinsic record references to adopt ETRI’s construction of RAID as requiring “a single logical unit,” while ignoring other extrinsic evidence that taught away from ETRI’s proposed construction. (See JA1502-3 (¶¶ 5-6)); *see Interval Licensing LLC v. AOL, Inc.*, 766 F.3d 1364, 1377 (Fed. Cir. 2014) (reversing narrower construction based on dictionary definitions in favor of broader definitions consistent with intrinsic evidence). *Helmsderfer*, 527 F.3d at 1382-83 (taking survey of dictionaries to find common disclosure among extrinsic evidence). But it is fundamental that “[a] claim should not rise or fall based on the preferences of a particular dictionary editor, or the court’s independent decision,

uninformed by the specification, to rely on one dictionary rather than another.” *Phillips*, 415 F.3d at 1322.

The Board’s determinations with respect to the extrinsic evidence are not supported by substantial evidence. Not one reference uses the term “a single logical unit.” Instead, the Board inexplicably equated, without any support, “single logical unit” with the phrase “single unit” in Weygant (JA759), the “high performance logical disk” described in Chen (JA1439),¹² the “single logical drive” described in Webster’s definition of some (but not all) of the RAID configurations (JA1776), and the “single storage unit” in the 2002 Microsoft dictionary definition (JA1779). (JA8-9). But both experts noted that “not all companies use all terms identically for technology” in this field, (*e.g.*, JA2073:8-12) and more specifically, they both testified that these various terms were imprecise and did not have identical meanings. (*E.g.*, JA2067:6-18 (ETRI’s expert testifying Webster’s Computer Dictionary definition is “a little imprecise”); JA20714-19 (same and noting 2002 Microsoft Computer Dictionary’s use of “single storage unit” imprecise); JA1502 (¶ 5) (Petitioners’ expert “note[d] that not all definitions of the term RAID require a single logical unit”)); *see Phillips*, 415 F.3d at 1322

¹² Furthermore, the cited section of Chen, which discusses a “large, high-performance disk,” is not referring to RAID; rather that section of Chen discusses certain types of disk arrays which are “highly vulnerable to disk failures,” which Chen explains should be overcome by introducing redundancy. (JA1439).

(“[D]ifferent dictionaries may contain somewhat different sets of definitions for the same words.”).

Even accepting that “single unit” and “single logical unit” are equivalent, as the Board contends, Weygant teaches that RAID is broader than a “single unit.” (JA759; *see also* JA 339-40; JA1502-03 at ¶ 6). In fact, the Board admits that Weygant’s definition of RAID includes not only “a single logical unit” “*but also* [] ‘various combinations of striped and mirrored configurations.’” (JA8) (emphasis added). Accordingly, the Board admits that Weygant discloses RAID configurations other than a single logical unit. (JA8 (“Weygant also indicates that a RAID can exist in forms other than a single logical unit,...”). Moreover, although the Board relied on Weygant’s disclosure of a “single unit,” ETRI’s expert, who agreed with the Board’s ultimate construction, testified that he did not “know how Weygant is using a single unit” and “that Weygant’s definition is nonsensical,” and otherwise was unable to define or describe what constitutes a “single logical unit.” (JA2079:23-25; JA2080:16-19; *see also* JA2077:3-10).

Where, as here, even the extrinsic evidence uses the term loosely and inconsistently, the term “single logical unit” should not have been imported into the construction of “RAID”. *See InterDigital*, 690 F.3d at 1327 (rejecting importation of term from the extrinsic record where extrinsic evidence demonstrated term was “used somewhat loosely by those working in the field of

cellular communications”). Without basis in either the extrinsic or intrinsic record, the Board’s narrowing of “RAID” to be limited to “a single logical unit” should be rejected in favor of the Board’s initial—and abundantly supported—construction of “RAID” as “redundant array of inexpensive disks.”

d. Neither the Intrinsic Nor Extrinsic Record Supports the Board’s Further Narrowing of “RAID” to Preclude Direct Communications With Storage Controllers.

The Board’s errors did not stop at limiting “RAID” to “a single logical unit.” The Board further limited “a single logical unit” to preclude the host computer from accessing storage controllers individually or directly. (JA23). But there is no discussion of any such requirement in any of the intrinsic *or* extrinsic evidence. (See JA339-40). Instead, as discussed in further detail in § III.C.2, *infra*, this narrowing is entirely driven by ETRI’s desire to exclude Hathorn’s disclosures from the construction of “RAID.” During the oral hearing, ETRI’s counsel argued that Hathorn’s “DASDs 326 and 336 are not together a single logical unit” because “the primary host 301 can individually access either one of those DASDs.” (JA417:4-7). The Board seized upon this attorney argument to limit “single logical unit” to explicitly exclude Hathorn’s teachings, stating “[w]e agree with [ETRI] that DASDs 326, 336 are not a single logical unit because primary host 301 *can* individually (or directly) access either one of those DASDs. (JA23 (citing only JA417:4-7)). Tellingly, it did not cite any intrinsic or extrinsic evidence in support

of such narrowing. That is because there is none. In fact, the intrinsic record—specifically, the claims—contemplates direct communication between the host and RAID controllers. (JA35 at 5:18-21).

This Court should rule that the broadest reasonable construction of “RAID” is “redundant array of inexpensive disks,” which a person of ordinary skill at the time would understand to inherently cover the known RAID configurations, including mirroring or shadowing (RAID Level 1), as discussed in the extrinsic evidence. In doing so, this Court should reject the Board’s construction of “RAID” as requiring “a single logical unit” that excludes RAID 1 mirroring or shadowing redundancy and excludes any direct communication between hosts and controllers within a RAID. The Board’s construction undermines the most fundamental principles of claim construction: it contradicts the intrinsic record, lacks basis in any extrinsic evidence, and improperly credits attorney argument.

C. Hathorn Anticipates the ’346 Patent Challenged Claims Under Both the Proper Construction and the Board’s Erroneous Construction of “RAID”.

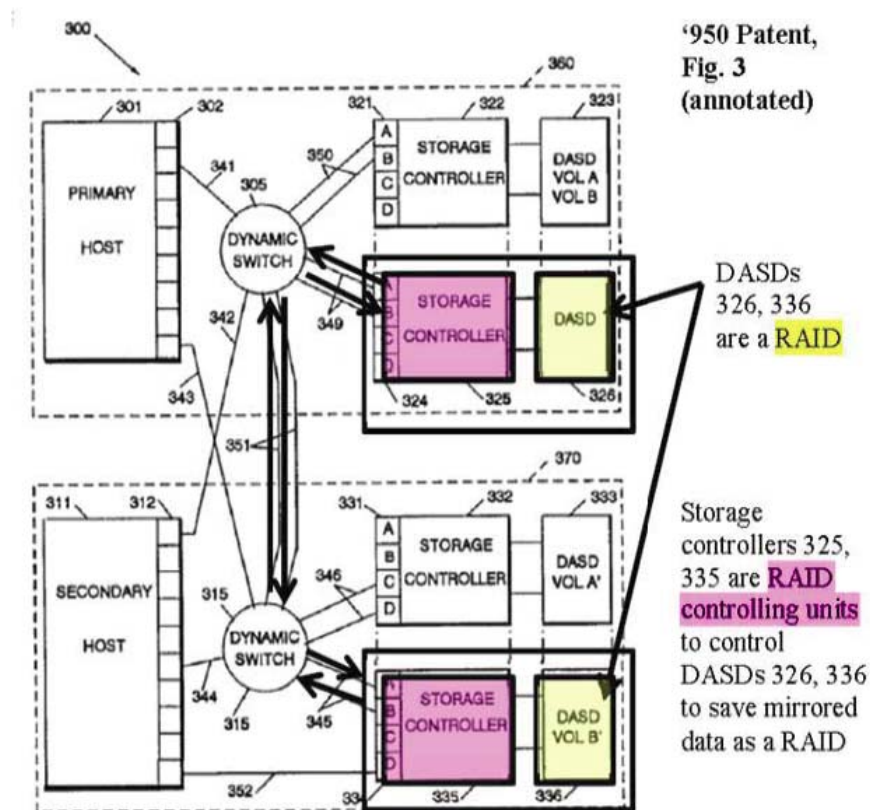
1. Hathorn Discloses a System with a Redundant Array of Inexpensive Disks In a Mirrored Configuration.

Hathorn discloses a “RAID” under the proper construction of the term—as a “redundant array of inexpensive disks.”

Hathorn describes back-up solutions such as “dual copy” configurations that include “mirroring” when “data is written to the additional DASDs [Direct Access

Storage Devices] ... [so that] if the primary DASDs fail, the secondary DASDs can be depended upon for data.” (JA1043 at 1:60-2:3). For example, “the primary data is stored on a primary DASD string attached to a primary storage controller while the secondary data is stored on a secondary DASD string attached to a secondary storage controller.” (JA1043 at 2:53-56).

Specifically, Figure 3 of Hathorn shows a data mirroring configuration which includes an array of multiple disks (DASDs) 323, 326, 333, and 336, and corresponding storage controllers 322, 325, 332 and 335, connected through dynamic switches 305 and 315 (as annotated):



(JA1528; JA335).

As shown in this figure, the “primary storage controllers [top purple box] ... send data or records for back-up directly to secondary storage controllers [bottom purple box].” (JA1046 at 7:32-44). Hathorn uses the term “shadowing” to describe the process of data being saved at DASD 326 (top yellow box) and then copied to DASD 336 (bottom yellow box) via storage controllers 325 and 335 (highlighted in purple). (JA1047 at 9:29-50 (“data written to the primary device [can] be shadowed to the secondary device” in Hathorn)). This copying of data to backup DASD 336 is what makes Hathorn’s disk array redundant. (JA85-86; JA194).

Persons of ordinary skill at the time of the invention understood that the shadowing and mirroring of data described in Hathorn (*e.g.*, between DASDs 326 and 336) is a RAID Level 1 configuration. (See JA20 (citing JA333)). For example, Chen, a preeminent publication on RAID technology, describes “Mirrored (RAID Level 1)” and states “[t]he traditional solution, called *mirroring* or *shadowing*, uses twice as many disks as a nonredundant array.” (JA1447 (italics in original; underlining added); JA1503 at ¶ 7; *see also* JA759 (“RAID 1: Disk Mirroring”); JA656 (RAID Level 1 is when “the controller writes data to [a] mirrored groups of disks”); JA335-36).

Indeed, the Board admitted that “RAID level 1, one of several possible RAID configurations, is ‘disk mirroring.’” (JA22 (citing JA1072)). The Board

further admitted that “Hathorn describes its ‘remote data shadowing’ as ‘mirroring,’” and “Hathorn may disclose a RAID level 1 configuration.” (JA22 (citing JA1072; JA2088:3-11; JA1043 at 1:65-67)).

Hathorn itself explains that its array of storage devices “DASDs can be configured as a RAID” and that “there are several different RAID configurations available.” (JA1043 at 2:5-11; JA337). The Board dismissed this express disclosure in Hathorn, asserting that “[b]eing capable of a RAID configuration is not the same as an actual RAID.” (JA22; *but see* JA196 (finding opposite at institution)). Without substantial analysis, the Board stated that it was unconvinced that “DASDs can be arranged as a RAID.” (JA22). The Board’s reasoning is erroneous in several respects. First, the Board cannot simply state it is unconvinced, when Hathorn expressly and undisputedly teaches that there are several different RAID configurations and that DASDs can be arranged as a RAID. (JA1043 at 2:5-11). Further, neither party’s expert disputed that DASDs can be configured as a RAID. (JA1876:4-18 (Dr. Mercer); JA1763-64 (¶ 47) (Dr. Conte)). Finally, the ’346 specification does not exclude from the scope of “RAID” those known prior art configurations of DASDs which constitute a RAID. *See supra* § III.B.1.a.

Because the Board did not rely on any other claim terms to distinguish the Hathorn reference (JA22-23), and because the Board’s Institution Decision

properly found the other limitations of the '346 patent claims to be satisfied by Hathorn (JA194-199), the Board's determination of no anticipation should be reversed. Thus, the Court should reverse the Board's Final Written Decision and find that the '346 patent claims at issue are anticipated by Hathorn.¹³

2. Hathorn Discloses a "RAID" Even Under the Board's Erroneous Construction

Even if this Court affirms the Board's construction of "RAID," the Court should still reverse the Board's determination of no anticipation because Hathorn discloses a RAID as construed to mean a "single logical unit."

As described in Hathorn's preferred embodiment, the host computer does not send the data to the secondary storage for back up; rather, the "primary storage controllers . . . send the data or records for back-up directly to secondary storage controllers." (JA1046:32-44). Thus, Hathorn discloses an embodiment in which the host computers communicate solely with the primary storage controllers of Hathorn's DASDs. By communicating only with the primary storage controllers, and teaching that these primary storage controllers in turn will communicate with the secondary controllers, Hathorn teaches that the host computers interact with the

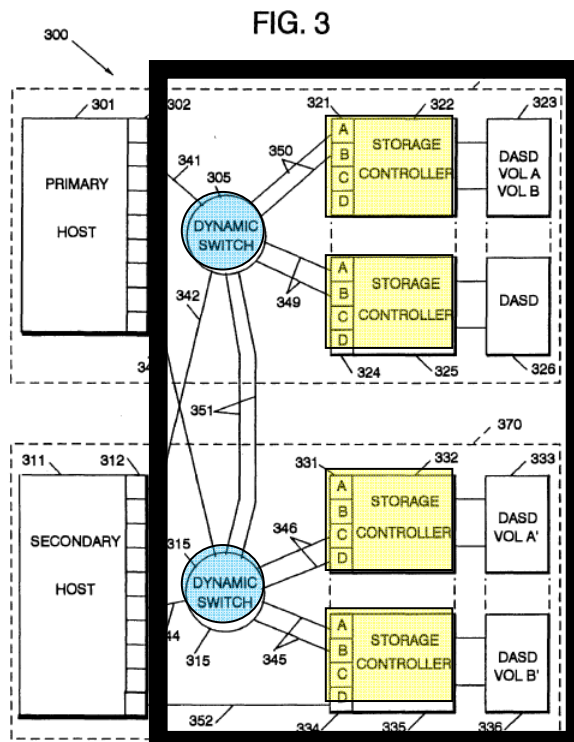
¹³ If the Court determines that the Board's construction is legally erroneous but adopts a construction of RAID other than "redundant array of inexpensive disks," Appellants respectfully request that this Court remand the IPR to the Board for further proceedings with respect to anticipation under the revised construction. *See Briantree Labs., Inc. v. Novel Labs., Inc.*, 749 F.3d 1349, 1352 (Fed. Cir. 2014); *TI Group Auto. Sys. (N.A.), Inc. v. VDO N.A., LLC*, 375 F.3d 1126, 1139-40 (Fed. Cir. 2004).

storage system as “a single logical unit,” as defined by the Board. (JA1047 at 8:6-9:51; *see also* JA2100:3-2101:2 (admitting Hathorn’s FIG. 3 could be viewed as disclosing a “single logical unit”)). Indeed, ETRI’s expert admitted that Hathorn’s storage controller performs the data shadowing operation of FIG. 6, (JA2098:19-24), and admitted that the primary host computer is not writing data directly to the secondary storage system. (JA2099:24-2100:2). Thus, Hathorn meets the “single logical unit” requirement of the Board’s construction of “RAID.”


In the proceedings below, ETRI referred to this “single logical unit” as a “black box.” (JA262 (quoting JA1745-46 (¶ 19)); JA404). In the Final Written Decision, the Board agreed that a RAID operating as a single logical unit would be such a black box. However, the Board ignored that ETRI’s expert admitted that Hathorn’s system, in a preferred embodiment, behaves as a “black box” as required by the “single logical unit” construction. (JA337-38 (quoting JA1768-69)). ETRI’s expert testified that in Hathorn’s Fig. 6 “*the primary host sees not only 325 and 335 as a single logical unit, but it also would see the combination of that plus 305, 351, 315, and 345, and of course the interface 335. That would constitute the black box that the primary host is interacting with.*” (JA2100:3-14 (emphases added)).

Thus, ETRI’s expert, Dr. Conte, admitted that Hathorn discloses a “single logical unit” that includes the storage controllers 325 and 335 (yellow), the

switches 305 and 315 (blue) with the corresponding links 351 and 345 and interfaces 334. (See FIG. 3 of Hathorn (reproduced below, with highlights and “black box” mark-up around the elements identified by Dr. Conte)).



Hathorn’s storage controllers (yellow), switches (blue), and corresponding links and interfaces are the circuits that make up a “RAID” in the ’346 patent. (Compare with JA30 (FIG. 4 of ’346 patent (reproduced below, with highlights and “black box” mark-up around “RAID 490”)). Thus, Hathorn discloses the “single logical unit” limitation of the “RAID” required by the Board.



The Board reasoned that Hathorn’s direct access storage devices “are not a single logical unit because primary host 301 *can* individually (or directly) access either one of those.” (JA23 (emphasis in original)). In support, the Board relied on ETRI counsel’s explanation that “the reason that the DASD 326 and 336 are not together a single logical unit is that the primary host 301 can individually access either one of those DASDs.” (JA23 (citing JA417:4-7)).

Citation to ETRI’s attorney argument is legal error and cannot substitute for the evidence in the record. The record shows that this “direct access” in Hathorn is just one deprecated embodiment—specifically, Hathorn describes the direct access cited by the Board as an embodiment which “may be undesirable due to the

required primary host resources” (JA1046 at 7:37-38).¹⁴ It is legal error for the Board to limit Hathorn to just this embodiment as a basis to distinguish the reference from the ’346 patent claims. *Arthrocare Corp. v. Smith & Nephew, Inc.*, 406 F.3d 1365, 1372 (Fed. Cir. 2005). The Board ignored that Hathorn expressly teaches that “[m]ore desirably” the primary storage controllers should send the data to the secondary storage controllers, which provides quicker communication. (JA1046 at 7:38-44). Thus, contrary to the Board’s reliance on ETRI’s counsel’s argument, Hathorn does not require that the hosts have individual or direct access to the corresponding storage devices. The Board’s narrow focus on one embodiment was reversible error. *See Hewlett-Packard Co. v. Mustek Sys., Inc.*, 340 F.3d 1314, 1324 n.6 (Fed. Cir. 2003) (“The anticipation analysis asks solely whether the prior art reference discloses and enables the claimed invention, and not how the prior art characterizes that disclosure or whether alternatives are also disclosed.”).

Moreover, the intrinsic evidence—in particular, claim 1’s language—expressly requires such access, providing that “the first RAID controlling unit and the second RAID controlling unit *directly exchange* information with the numerous host computers....” (JA35 at 5:18-21(emphasis added); *see* JA338-40). Thus,

¹⁴ The Board implicitly admits that this embodiment is optional when it states “DASDs 326, 336 are not a single logical unit because the primary host *can*”—not must—“individually or directly access either of the DASDs.” (JA23 (emphasis in original)).

narrowing “RAID” to a “single logical unit” is not only unfounded in the extrinsic record, it plainly conflicts with the intrinsic record. This manufactured limitation should not and cannot preclude Hathorn from anticipating claims 1-3 and 5-8 of the ’346 patent.

IV. CONCLUSION

For the foregoing reasons, this Court should reverse the Board’s Final Written Decision.

Date: November 23, 2015

Respectfully submitted,

By: /s/ Peter H. Kang

Thomas M. Dunham
James M. Woods
Winston & Strawn LLP
1700 K Street, NW
Washington, DC 20006
Telephone: (202) 282-5863

Counsel for Appellant Dell Inc.

Peter H. Kang
SIDLEY AUSTIN LLP
1001 Page Mill Road, Building 1
Palo Alto, CA 94304
Telephone: (650) 565-7006

Anna M. Weinberg
SIDLEY AUSTIN LLP
1501 K Street, N.W.
Washington, D.C. 20005
Telephone: (202) 736-8000

*Counsel for Appellants Dell Inc.,
Hewlett Packard Enterprise Company,
and NetApp, Inc.*

CERTIFICATE OF COMPLIANCE

This brief complies with the type-volume limitations of Federal Rule of Appellate Procedure 32(a)(7)(B) and the Rules of this Court, because it contains 12,992 words (as determined by the Microsoft Word 2007 word-processing system used to prepare the brief), excluding the parts of the brief exempted by Federal Rule of Appellate Procedure 32(a)(7)(B)(iii).

This brief complies with the typeface requirements of Federal Rule of Appellate Procedure 32(a)(5) and the type-style requirements of Federal Rule of Appellate Procedure 32(a)(6) because it has been prepared in a proportionally spaced typeface using the Microsoft Word 2007 word-processing system in 14-point Times New Roman font.

/s/ Peter H. Kang

Peter H. Kang.

CERTIFICATE OF SERVICE

I hereby certify that on November 23, 2015, a true and correct copy of the foregoing was timely filed with the Clerk of the Court using the appellate CM/ECF system, which will send notifications to all counsel registered to receive electronic notices.

/s/ Peter H. Kang
Peter H. Kang

ADDENDUM

ADDENDUM TABLE OF CONTENTS

Final Written Decision (02/27/2015).....	A1
U.S. Patent 6,978,346.....	A26

Trials@uspto.gov

571-272-7822

Paper 39

Entered: February 27, 2015

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

DELL INC., HEWLETT-PACKARD COMPANY,
and NETAPP, INC.,
Petitioners,

v.

ELECTRONICS AND TELECOMMUNICATIONS
RESEARCH INSTITUTE,
Patent Owner.

Case IPR2013-00635
Patent 6,978,346 B2

Before BRIAN J. McNAMARA, MIRIAM L. QUINN, and
GREGG I. ANDERSON, *Administrative Patent Judges*.

ANDERSON, *Administrative Patent Judge*.

FINAL WRITTEN DECISION
35 U.S.C. § 318(a) and 37 C.F.R. § 42.73

IPR2013-00635

Patent 6,978,346 B2

I. INTRODUCTION

On September 27, 2013, Dell, Inc., Hewlett-Packard Company, and NETAPP, Inc. (collectively, “Petitioner”) filed a Petition requesting an *inter partes* review of claims 1 through 9 of U.S. Patent No. 6,978,346 B2 (Ex. 1001, “the ’346 patent”). Paper 1 (“Pet.”). On March 20, 2014, we instituted trial for claims 1–3 and 5–8 of the ’346 patent on certain of the grounds of unpatentability alleged in the Petition. Paper 19 (“Decision on Institution” or “Dec. Inst.”).

After institution of trial, Electronics and Telecommunications Research Institute (“Patent Owner”) filed a Patent Owner Response. Paper 28 (“PO Resp.”). Petitioner filed a Reply. Paper 33 (“Pet. Reply”).

An oral hearing was held on December 18, 2014. The transcript of the consolidated hearing has been entered into the record. Paper 38 (“Tr.”).

We have jurisdiction under 35 U.S.C. § 6(c). This final written decision is issued pursuant to 35 U.S.C. § 318(a).

A. *Related Proceedings*

The ’346 patent has been asserted against Petitioner in the following actions: *Safe Storage LLC v. Dell Inc.*, 1-12-cv-01624 and *Safe Storage LLC v. NetApp Inc.*, 1-12-cv-01628. Pet 1–2. Petitioner advises us of an additional seventeen actions involving the ’346 patent against third parties, all pending in the United States District Court for the District of Delaware. *Id.*

B. *The ’346 Patent*

The ’346 patent describes an apparatus with “redundant interconnection between multiple hosts and a redundant array of inexpensive disks (hereinafter referred to as ‘RAID’).” Ex. 1001, Abstract. As a result

IPR2013-00635

Patent 6,978,346 B2

of the redundant interconnection, the apparatus allows increased bandwidth in the event one of the two RAID controllers 460 and 461 has a failure. *Id.* at 3:1–9.

Figure 4 of the '346 patent is reproduced below:




Figure 4 is a block diagram of a host matching system including RAID 490 and its interconnection to host computers 400–405. Ex. 1001, 2:64–3:6. RAID 490 includes two RAID controllers 460, 461 and hubs 440, 441. *Id.* at 3:10–18. Each RAID controller includes a pair of network interface controllers. For example, RAID controller 460 includes network interface controllers 470, 471, and RAID controller 461 includes network interface controllers 480, 481. *Id.* at 3:11–13. Each host computer has its own network interface controller (410–415), which connects the host computer through the hubs and to the network interface controllers (470, 471, 480, 481) of RAID controllers 460, 461. *Id.* at 3:31–35.

IPR2013-00635

Patent 6,978,346 B2

The '346 patent describes that the result is two independent networks with twice the bandwidth of a single network and a “communication passage” between the two RAID controllers. *Id.* at 3:62–64. The communication passage creates a “fault tolerant function” should one of the RAID controllers 460 or 461 fail. *Id.* at 3:64–66. According to Figure 4, communications line 450 interconnects network interface controller 480 of RAID controller 461 and network interface controller 470 of RAID controller 460. *Id.* at 4:2–6; Fig. 4. Then, RAID controller 461 may send information to RAID controller 460. *Id.* In like manner, network interface controller 471 of RAID controller 460 may be connected over communications lines to network interface controller 481 of RAID controller 461, allowing RAID controller 460 to send information to RAID controller 461. *Id.* at 3:66–4:2.

In summary, and as shown in Figure 4, a communication circuit is provided for an error recovery, while maintaining bandwidth communication between two RAID controllers 460, 461. Ex. 1001, 3:1–5. Even though one RAID controller 460 or 461 has an occurrence of a trouble, the bandwidth becomes twice the single connection bandwidth. *Id.* at 3: 6–9.

C. Illustrative Claim

Independent claim 1 is reproduced below:

1. An apparatus for a redundant interconnection between multiple hosts and a RAID, comprising:
 - a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers, the first RAID controlling unit including a first network controlling unit and a second network controlling unit, and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit; and

IPR2013-00635

Patent 6,978,346 B2

a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers, wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the numerous host computers through the plurality of connecting units, and the first network controlling unit exchanges information with the fourth network controlling unit, and the second network controlling unit exchanges information with the third network controlling unit.

D. Ground Upon Which Trial Was Instituted

Trial was instituted on the ground alleging that claims 1–3 and 5–8 of the '346 patent are anticipated under 35 U.S.C. § 102(b) by Hathorn, U.S. Patent No. 5,574,950, issued November 12, 1996.

II. ANALYSIS

A. Claim Construction

1. Principles of Law

In an *inter partes* review, claim terms in an unexpired patent are interpreted according to their broadest reasonable construction in light of the specification of the patent in which they appear. 37 C.F.R. § 42.100(b).

2. “RAID” (Claim 1)

In the Decision on Institution we found that “RAID” is well understood by a person of ordinary skill in the art as an acronym for “redundant array of inexpensive disks.” Dec. Inst. 8 (citing Ex. 1001, Abstract). Patent Owner does not dispute the interpretation, but points out that each word of the construction conveys additional significance. PO Resp. 10.

With regard to the word “disks,” Patent Owner argues that “disks” means “disk drives,” and that a RAID is an “array of multiple disk drives configured for redundancy.” *Id.* (citing Declaration of Dr. Thomas M.

IPR2013-00635

Patent 6,978,346 B2

Conte, Ex. 2003 ¶¶ 16, 18). Based on the presence of “array” in our preliminary construction, Patent Owner offers evidence that an “array” is “a single logical storage unit of disk drives.” *Id.* at 11 (citing *Webster’s Computer Dictionary* 308 (9th ed. 2001) (Ex. 2004, 11); *Microsoft Computer Dictionary* 437 (5th ed. 2002) (Ex. 2005, 3)¹; Ex. 2003 ¶ 19). Patent Owner also directs us to Dr. Conte’s testimony for additional support. *See* Ex. 2003 ¶¶ 16–21. Patent Owner also cites to Weygant, which it alleges “only states that combinations of striping and mirroring do not appear as a single logical unit, but other forms of RAID like RAID Level 1 mirroring do.” Tr. 38:9–13; Ex. 1003, 153.² Lastly, Patent Owner points to Chen which states that disk arrays “organize multiple independent disks into a large high-performance logical disk.” Tr. 38:14–19; Ex. 1011, 5. Thus, Patent Owner contends that “RAID” should be construed as “a single logical unit for mass storage that provides fault tolerance and recovery via employing multiple physical disk drives.” PO Resp. 12 (citing Ex. 2003 ¶ 38).

Petitioner proposes a construction of “RAID” consistent with the construction we provided in the Decision on Institution. Pet. 8. In response to Patent Owner’s proposal, Petitioner argues that “the data mirroring

¹ The dictionary definitions are dated in 2001 and 2002, respectively, after the foreign priority date of the ’346 patent, September 19, 2000. *See* Exs. 2004 and 2005. However, references having publication dates after the critical date may be cited to show the state of the art at or around the time of the invention. *Eli Lilly and Co. v. Barr Labs., Inc.*, 251 F.3d 955, 969–70 (Fed. Cir. 2001); *see also In re Wilson*, 311 F.2d 266, 268–269 (CCPA 1962).

² Citations to Weygant (Ex. 1003) are to Weygant’s pages and not the Exhibit page number.

IPR2013-00635

Patent 6,978,346 B2

operation of Hathorn is a RAID configuration.” Pet. Reply 1.³ Petitioner offers evidence that data mirroring is RAID Level 1. *Id.* at 3–4 (citing Declaration of Dr. M. Ray Mercer, Ex. 1006, 22–23; Second Declaration of Dr. M. Ray Mercer, Ex. 1012, 3–4). Both cited portions of Dr. Mercer’s Declarations cite to Weygant (Exhibit 1003) as supporting Dr. Mercer’s opinion that “RAID” may be construed differently depending upon the particular RAID configuration, i.e., RAID Level 1–5. Pet. Reply at 7–8; Ex. 1003, 153. Dr. Mercer also cites to a 1999 edition of *Microsoft Computer Dictionary*, where the definition of “RAID” does not include “single logical storage unit.” Ex. 1006, 22. Dr. Mercer testifies that, at the time the application for the ’346 patent was filed, on December 29, 2000, “there is no one definition of the term that is agreed by everyone.”⁴ *Id.* at 21. Relying on Weygant, Dr. Mercer concludes that, in a RAID, “a group of disks do[es] not have to be configured as a single unit.” *Id.* at 4.

We first review the intrinsic evidence as to the meaning of “RAID.” The written description of the ’346 patent restates the acronym for RAID, but otherwise lacks additional description of RAID or its functionality. Consistently and throughout the written description, RAID is referred to in the singular, i.e., “the apparatus for a redundant interconnection between multiple hosts and *a RAID* comprises a plurality of RAID controllers.” Ex.

³ Petitioner’s Reply Brief lacks page numbers. Page number references used here begin with page 1, the first page following page iv.

⁴ Dr. Mercer cites to the 1999 edition of *Microsoft Computer Dictionary*, which is not an exhibit of record, and, therefore, constitutes inadmissible hearsay, absent an exception. Nevertheless, the excerpt relied on need not be admissible for the opinion of Dr. Mercer regarding the definition of “RAID” to be admitted. Fed. R. Evid. 703. We further note that Patent Owner does not allege that the excerpt should be excluded.

IPR2013-00635

Patent 6,978,346 B2

1001, 2:16–18 (emphasis added). The claims also recite “a RAID.” *Id.* at 5:7–8. Figure 4 of the ’346 patent shows RAID 490 as a single component within a box which includes two RAID controllers 460 and 461. Similarly, the ’346 patent represents the prior art RAID as a single component. Ex. 1001, Fig. 1, element 130; Fig. 2, element 240; Fig. 3, element 340. Neither party relies on the prosecution history (Exhibit 1002), and our independent review of that history failed to reveal any additional insight as to the term’s meaning.

We now turn to the extrinsic evidence. Patent Owner’s expert, Dr. Conte, testifies that a RAID is a mass storage device built from multiple, physical disk drives. Ex. 2003, 9. This evidence is uncontroverted. The prior art supports Dr. Conte. Weygant discloses that a RAID is a single logical unit, but also in “various combinations of striped and mirrored configurations.” Ex. 1003, 153. Chen defines RAID to be Redundant Arrays of Inexpensive *Disk*s. Ex. 1011, 1. (emphases added). Chen’s discussion of RAID technology states that the problem of obtaining high performance is addressed by “arrays, which organize multiple independent disks into *a large, high-performance logical disk.*” *Id.* at 2 (emphasis added).

The Specification consistently refers to RAID in the singular. Both Weygant and Chen, which predate the effective filing date of the ’346 patent, and the testimony of Dr. Conte, support that a RAID is a single logical unit. Although Weygant also indicates that a RAID can exist in forms other than a single logical unit, those forms are limited to “striped and mirrored” configurations. Ex. 1003, 153. Furthermore, although the two dictionary definitions cited by Patent Owner are found in dictionaries

IPR2013-00635

Patent 6,978,346 B2

published after the foreign priority date of the '346 patent, they further substantiate the proposition that those of ordinary skill generally refer to RAID as a single logical unit. Exs. 2004, 2005. We have not been shown evidence that the RAID of the '346 patent is configured in a “striped and mirrored configuration,” which might not be in the single unit configuration of RAID. *See* Ex. 1003, 153. The evidence, on the full record before us, persuades us that it is necessary to modify our preliminary construction of “RAID” to account for the understanding of a person of ordinary skill in the art in light of the '346 patent disclosure. Accordingly, applying the broadest reasonable interpretation we construe “RAID,” as the term is used in the '346 patent, to mean “a single logical unit for mass storage using multiple physical disk drives.”

3. *“RAID controlling unit” (Claim 1)*

In the Decision on Institution we found “RAID controlling unit” to mean “a component that controls operation of the RAID.” Dec. Inst. 9–10. In our analysis, we declined to include extraneous language unsupported by either the '346 patent or extrinsic evidence. *Id.* Patent Owner states that our interpretation “is not incorrect,” but points out that each word of the construction conveys additional significance. PO Resp. 12. Patent Owner argues that the function of a RAID controller is to provide redundancy by writing redundant data to multiple disk drives. *Id.* Thus, either a RAID controller, or multiple RAID controllers, “must be able to write to all of the disk drives in the RAID unit in order to perform redundancy.” *Id.* at 12–13 (citing Ex. 2003 ¶¶ 39, 41). Patent Owner contends that there must be a second RAID controller to establish the redundancy it argues is required for a RAID. PO Resp. 42–46. Patent Owner proposes that “RAID controlling

IPR2013-00635

Patent 6,978,346 B2

unit” be construed as “a component that controls operation of the RAID so as to provide redundant storage of data among the array of disk drives.” *Id.* at 13.

Petitioner argues that the construction of “RAID controlling unit” from the Decision on Institution should not be further narrowed. Pet. Reply 8. Also, Petitioner disagrees with Patent Owner that “a RAID controlling unit must *directly* access *all* disks in the RAID.” *Id.*

We find that Patent Owner’s proposal imports limitations not supported by the Specification, e.g., “among the array of disk drives.” Other than its appearance in the Abstract of the ’346 patent, as part of the acronym for RAID, “array” does not appear in the Specification of the ’346 patent. We decline to go beyond the ’346 patent to add limitations to the construction not supported by the Specification, particularly when Patent Owner’s expert agrees with our preliminary construction. Ex. 2007, 64:7–11.

That we changed our construction of RAID does not impact our construction of “RAID controlling unit.” We are presented with no compelling reason to change our construction from the Decision on Institution. Accordingly, applying the broadest reasonable interpretation to the term RAID controlling unit, we construe “RAID controlling unit” to mean “a component that controls operation of the RAID.”

4. “*First RAID controlling unit*” and
“*Second RAID controlling unit*” (Claim 1)

“First RAID controlling unit” and “second RAID controlling unit” were not interpreted in the Decision on Institution, but Patent Owner argues

IPR2013-00635

Patent 6,978,346 B2

the terms should be construed. PO Resp. 13. Based on the claim language and Specification, Patent Owner argues that the “first” and “second RAID controlling units” are for “the same RAID.” *Id.* Patent Owner asserts the only RAID in the claim appears in the preamble and should be given weight. *Id.* at 14 (citing *Proveris Sci. Corp. v. Innovasystems, Inc.*, 739 F.3d 1367, 1372 (Fed. Cir. 2014) (citations omitted)). Petitioner does not specifically argue against Patent Owner’s construction and has no proposal of its own.

As discussed above, we have construed RAID controlling unit. Claim 1 recites “a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers.” The Specification states that the RAID supports “a fault tolerance of RAID controllers and simultaneously heightening a performance.” Ex. 1001, Abstract, 1:12–14. Further, “a RAID comprises a plurality of RAID controllers for processing requests of numerous host computers connected with one another.” *Id.* at 2:17–19. We agree that the claim and disclosure both include “a RAID” and two or more RAID controlling units. This is supported by the claim language, which recites a RAID and two RAID controlling units. The claim does not recite an express numerical correspondence between a RAID and the RAID controlling units. Nevertheless, we have construed “RAID” and “RAID controlling unit,” above, and are satisfied that no construction of additional similar terms (“first RAID controlling unit” and “second RAID controlling unit”) is needed.

5. “*exchange/exchanges information*” (Claim 1)

Neither party argues the “exchange/exchanges information” limitation of claim 1. As we found in the Decision on Institution, claim 1 uses

IPR2013-00635

Patent 6,978,346 B2

“exchange” and “exchanges information” according to their ordinary sense: to transmit and receive information reciprocally.⁵ The claim recites the structures between which information is exchanged, i.e., between the RAID controlling units and the host computers, between the first and fourth network controlling units, and between the second and third network controlling units. The claim language requires only the information to and from the host computers to be exchanged through the connection units. The specification of the ’346 patent is consistent with the ordinary meaning of giving and receiving information reciprocally, because it describes that information is transmitted to and from a network interface controller of a first RAID and another network interface controller of a second RAID. Ex. 1001, 3:66–4:12.

We are presented with no reason to change our construction from the Decision on Institution. Accordingly, applying the broadest reasonable construction in light of the Specification, we construe “exchange/exchanges information” to mean “to transmit and receive information reciprocally.” Dec. Inst. 10–11.

6. “*network controlling unit*”

“Network controlling unit” was not interpreted in the Decision on Institution, but Patent Owner argues the term should be construed and that the term is “generally understood to one skilled in the art as a hardware controller that supplies communication functionality when attached to a computer network.” PO Resp. 16 (quoting Ex. 2003 ¶ 42). Patent Owner

⁵ *Definition exchange (vb) (3)*, WEBSTER’S THIRD NEW INTERNATIONAL DICTIONARY, UNABRIDGED (1993), *available at* <http://lionreference.chadwyck.com> (Dictionaries/Webster’s Dictionary).

IPR2013-00635

Patent 6,978,346 B2

submits that the '346 patent discloses two ports for each controller, one for transmitting and one for receiving. *Id.* Patent Owner, thus, proposes that the term additionally “includ[es] one or more ports.” *Id.* at 16–17. A network controlling unit having one port, according to Patent Owner, is too limiting and not disclosed in the Specification. *Id.*

Petitioner argues that the term should be given the plain and ordinary meaning thereof.⁶ Pet. Reply 9. In support for its argument, Petitioner reminds us that we disagreed in the Decision on Institution with the assertion in the Petition that a “network controlling unit” is necessarily specific hardware. *Id.* Petitioner points out that the Specification does not mention ports in connection with the “network controlling units,” and only in connection with the “hub” or “switch.” *Id.* at 10. Through its expert, Petitioner proposes a construction: “any component allowing a device to communicate over a network.” *Id.* (citing Ex. 1012, 5–8). Petitioner contends that its proposal is consistent with the plain and ordinary meaning of the term. *Id.*

The experts for both parties agree that “network controlling unit” refers to a device for communication connected to a network. We do not agree with Patent Owner that the broadest reasonable interpretation requires that “one or more” ports be part of the construction. PO Resp. 16–17. However, there must be a connection to the network for communication to occur. Accordingly, we construe “network controlling unit” as “a

⁶ Petitioner appears to rely on the Decision on Institution where we said that “Other Terms for Proposed Construction” would be given their ordinary and customary meaning and would not be construed “at this time.” Dec. Inst. 11.

IPR2013-00635

Patent 6,978,346 B2

component, connected to a network, for providing communication over the network.”

7. “*connection unit/hub/switch*” (Claim 5)

In the Decision on Institution, we found that, consistent with the definition provided in the Specification, “connection unit” is “a hub or switch.” Ex. 1001, 3:13–18. The preceding was Petitioner’s proposal in the Petition. Pet. 6. Patent Owner acknowledges the preceding definition but argues that “hub” and “switch” are not one in the same. PO Resp. 17. Patent Owner has no specific proposal for either “hub” or “switch” and only argues they are different in some unspecified way. Neither of the parties’ experts testifies as to any difference between these terms.

We find that the Specification treats “hub” and “switch” as equivalents. For example, Figure 4 shows components 440 and 441 labeled as a “HUB OR SWITCH.” We are not persuaded by Patent Owner’s argument that “hub” should be interpreted to exclude or be different from a “switch.” PO Resp. 19.

B. Anticipation By Hathorn

Petitioner contends that claims 1–3 and 5–8 of the ’346 patent are anticipated under 35 U.S.C. § 102(b) by Hathorn. Pet. 9, 45–60. To support this position, Petitioner relies on the testimony of Dr. Mercer. Ex. 1006, 130–163 (including claim chart). For the reasons discussed below, we determine that Petitioner has not shown by a preponderance of the evidence that claims 1–3 and 5–8 are unpatentable as anticipated by Hathorn.

1. Hathorn Overview

Hathorn discloses a remote copy system with dynamically modifiable ports on the storage controller that are alternatively configurable. Ex. 1005,

IPR2013-00635

Patent 6,978,346 B2

Abstract. A primary storage controller can appear as a host processor to a secondary storage controller. *Id.* Hathorn describes a method for communicating between host processors and storage controllers, or between storage controllers.

Figure 3 of Hathorn is reproduced below.

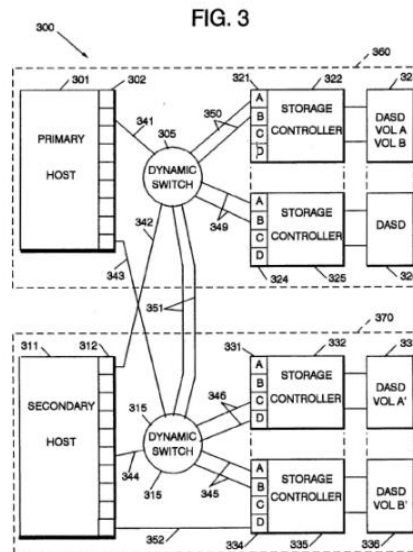


Figure 3 is a block diagram of one embodiment of a remote dual copy system of the invention described in Hathorn. Primary storage controller 322 communicates through port A 321 with secondary storage controller 332. Ex. 1005, 8:11–15. As shown in Figure 3, port A 321 acts as a channel link-level facility through communication links 350, dynamic switch 305, communication links 351, dynamic switch 315, and communication links 346 to communicate with secondary storage controllers 332 and/or 335. *Id.*

2. Analysis

Our discussion focuses on claim 1, the only independent claim at issue. All remaining challenged claims depend, directly or indirectly, from claim 1.

IPR2013-00635

Patent 6,978,346 B2

Petitioner cites to Hathorn's primary and secondary hosts and two dynamic switches 305, 315 as redundant interconnections recited in the preamble of claim 1. Pet. 46–47 (citing Ex. 1005, Fig. 3; Ex. 1006, 132–134). Petitioner also points out that Hathorn describes a RAID configuration that can be used in connection with a direct access storage device (“DASD”). *Id.* (citing Ex. 1005, 2:5–11).

Petitioner further points to Hathorn's storage controllers 322, 325, and 332, 335, respectively, as RAID controlling units that process requests from the primary host and secondary host for transferring data or records from the DASDs. Pet. 48 (citing Ex. 1006, 135–136). Petitioner also alleges that: (1) the communication ports in the storage controllers meet the “network controlling unit” limitation; (2) the first and second network controlling units are met by ports A, B 324; and (3) the third and fourth network controlling units are met by ports A, B 334. Pet. 48–49 (citing Ex. 1005, 8:5–6; Ex. 1006, 137–139).

Petitioner alleges that Hathorn's dynamic switches 305, 315 meet the recited “connection units.” Pet. 50 (citing Ex. 1005, Fig. 3; Ex. 1006, 139–141). The switches connect to the RAID controlling units by links 351. *Id.*

Hathorn discloses the following concerning Figure 3:

primary storage controller 322, via port A 321, can communicate with primary host 301 by communication links 350, dynamic switch 305 and communication link 341, wherein port A 321 is a control unit link-level facility. Alternately, primary storage controller 322, via the same port A 321, can communicate with secondary storage controller 332 by communication links 350, dynamic switch 305, communication links 351, dynamic switch 315, and communication links 346, wherein port A 321 acts as a channel link-level facility.

IPR2013-00635

Patent 6,978,346 B2

Ex. 1005, 8:6–15. Petitioner alleges the preceding disclosure and other similar disclosures in Hathorn disclose that the RAID controlling units “exchange information” through the connection units as claimed. Pet. 51–52 (citing Ex. 1005, 7:28–35, 8:3–15, Fig. 6, step 601; Ex. 1006, 141–142).

Petitioner contends that Hathorn explains that ports A and B 334, i.e., the third and fourth network controlling units, initiate the operation of Figure 4. Pet. 52–53 (citing Ex. 1005, 8:3–15; Ex. 1006, 143–144). Similarly, Petitioner argues that ports A and B 324, i.e., the first and second network controlling units, perform the data mirroring of Figure 5. *Id.* (citing Ex. 1005, 8:61–63, 9:49–51; Ex. 1006, 143–144). Petitioner relies on the preceding evidence to support that Hathorn discloses the second “exchanges information” limitation of claim 1.

Patent Owner relies on its construction of RAID and argues that “Hathorn also supports the view that a RAID ‘array’ must be a single logical storage unit of disk drives.” PO Resp. 11. Patent Owner contends that Hathorn:

makes a clear distinction between a “RAID” and a mirroring or dual-copy system employing two disk drives, which, although redundant, do not form an ‘array’ in the sense of a RAID. Indeed, Hathorn describes these two scenarios as “alternative[s].”

PO Resp. 11 (citing Ex. 1005, 1:60–2:11).

Beyond arguing that Hathorn does not show a RAID according to Patent Owner’s proposed construction, Patent Owner emphasizes the difference between Hathorn’s DASDs and a RAID. PO Resp. 22–23. Specifically, Patent Owner points to Petitioner’s reliance on Hathorn’s

IPR2013-00635

Patent 6,978,346 B2

Figure 3 and notes that Figure 3 makes no mention of “RAID,” only DASDs. *Id.*

Patent Owner emphasizes that Petitioner’s reliance on the sole mention of RAID in Hathorn is made in the Background discussion at column 2, lines 4 through 11. PO Resp. 23–24. The pertinent portion of Hathorn is as follows:

Another data back-up alternative that overcomes the need to double the storage devices involves writing data to a redundant array of inexpensive devices (RAID) configuration. In this instance, the data is written such that the data is apportioned amongst many DASDs. If a single DASD fails, then the lost data can be recovered by using the remaining data and error correction procedures. Currently there are several different RAID configurations available.

Ex. 1005, 2:4–11. Patent Owner argues that the cited disclosure from Hathorn distinguishes a RAID configuration from the “single DASD,” concluding that Hathorn discloses that a RAID would “be formed of ‘many DASDs.’” PO Resp. 23–24(citing Ex. 1005, 2:8). Again, Patent Owner points out that there is no mention of Figure 3 in the cited disclosure. *Id.* at 24.

Patent Owner disagrees with the Decision on Institution’s statement that “[t]he use of a RAID is disclosed specifically in Hathorn as a type of DASD.” PO Resp. 24 (citing Dec. Inst. 15). Petitioner relies on its expert, Dr. Conte, for support, citing the following: (1) Figure 3 of Hathorn does not mention RAID; (2) the column 2 discussion in Hathorn is not connected to Figure 3; and (3) Hathorn states that a RAID is made up from “many DASDs.” *Id.* at 25 (citing Ex. 2003 ¶ 52).

IPR2013-00635

Patent 6,978,346 B2

Patent Owner also argues that it is improper to combine the column 2 disclosure of Hathorn with Figure 3 to find that claim 1 is anticipated. PO Resp. 28. Thus, Hathorn does not disclose the limitations of claim 1 “arranged or combined in the same way as recited in the claim, [and] it cannot be said to prove prior invention of the thing claimed, and cannot anticipate under 35 U.S.C. § 102.” *Id.* at 27 (citing *Net MoneyIN, Inc. v. Verisign, Inc.*, 545 F.3d 1359, 1371 (Fed. Cir. 2008)). Patent Owner also argues that combining the Background description in column 2 with the embodiment of Figure 3 of Hathorn “is in the province of an obviousness inquiry, not anticipation.” *Id.* at 28. The ground at issue here is anticipation, and Patent Owner concludes that the challenge must fail. *Id.*

Patent Owner contends that Petitioner clarified its position on how Hathorn shows a RAID through the deposition of Petitioner’s expert, Dr. Mercer. PO Resp. 34–35. Patent Owner characterizes the Petitioner’s challenge as “a mirrored or remote-copy pair of different DASDs in Figure 3 constitute a RAID.” *Id.* (citing Ex. 2006, 97⁷). Patent Owner argues two DASDs do not form a RAID because they do “not form a single logical unit or drive.” *Id.* at 36 (citing Ex. 2003 ¶ 59).

In addition to its argument that Hathorn’s Figure 3 does not disclose a RAID configuration, Patent Owner argues Hathorn lacks “first” and “second RAID controlling units.” PO Resp. 37–48. Patent Owner also argues that Hathorn fails to disclose a “RAID controlling unit” with two “network controlling units.” *Id.* at 48–50. Patent

⁷ See Ex. 2006, 97:4–18.

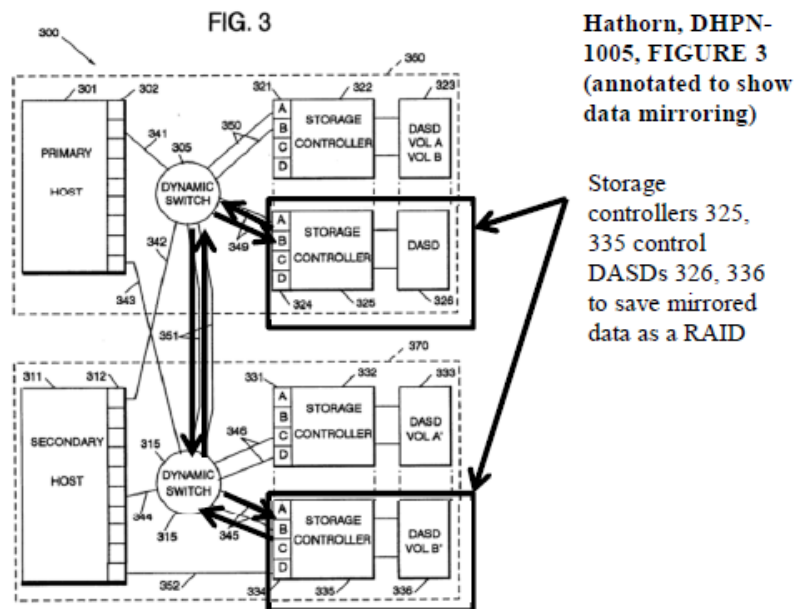
IPR2013-00635

Patent 6,978,346 B2

Owner argues that Hathorn fails to anticipate claims 2, 3, and 8 under the Petition's interpretation of "coupled" and "connected." *Id.* at 50–54. Turning to claim 5, Patent Owner argues Hathorn fails to disclose a "hub" as claimed. *Id.* at 58. Patent Owner's last argument relates to connection of connection ports with the host computer, as recited in claims 5–7. *Id.* at 55–57.

Petitioner responds that the data mirroring operation of Hathorn is a RAID configuration. Pet. Reply 1. And Petitioner also contests Patent Owner's arguments regarding the "RAID controller," arguing that Hathorn's storage controllers satisfy the broadest reasonable interpretation of RAID controller, i.e., "a component that controls operation of the RAID." *Id.* at 2.

Petitioner argues that Figure 3 of Hathorn shows data being shadowed or mirrored between DASDs 326, 336 through storage controllers 325, 336. Pet. Reply 3. Petitioner's annotation of Hathorn Figure 3 is reproduced below.



IPR2013-00635

Patent 6,978,346 B2

Figure 3, according to Petitioner, shows how Hathorn mirrors data between DASD 326 and DASD 336. *Id.* Petitioner argues that disk mirroring is RAID level 1. *Id.* at 3–4 (citing Ex. 1006, 22–23, 142–144; Ex. 1012, 3–4).⁸ Petitioner contends that Hathorn’s two storage controllers are two RAID controlling units because each provides control for the mirroring operation. *Id.* at 4–5. Further, Petitioner argues DASDs 326, 336 make up a RAID. Tr. 11:1–6.

Petitioner concludes that, because data mirroring is RAID level 1, Hathorn discloses a RAID. Pet. Reply 5. Petitioner’s argument continues to rely on the Background disclosure in Hathorn, which references a “RAID.” *Id.* Petitioner also notes that Patent Owner’s expert, Dr. Conte, acknowledges that Hathorn’s mirroring involving two storage controllers 325, 335 and two switches 305, 315 “hypothetically would constitute a RAID controlling Unit A.” *Id.* (citing Ex. 2003, 32–33).

In response to Patent Owner’s argument that a RAID controlling unit must be able to write directly to all disks in a RAID, Petitioner argues that this feature, which may be advantageous, is not recited in the claims. Pet. Reply 6. Further, Petitioner argues the claims express no relationship between the RAID and RAID controlling units. *Id.*

Petitioner has additional arguments relating to other limitations of claim 1 and to the dependent claims. *See* Pet. Reply 8–15.

Because our decision rests on the analysis of “RAID” and “RAID

⁸ Exhibits 1006 and 1012, Dr. Mercer’s first and second declarations, in turn cite to Weygant and Chen, Exhibits 1003 and 1011, for additional support that data mirroring is RAID level 1.

IPR2013-00635

Patent 6,978,346 B2

controlling unit” limitations, we need not repeat those arguments in detail here.

Petitioner cites to the Background section of Hathorn to establish that “DASDs can be *configured* as a RAID.” Pet. Reply 5 (citing Ex. 1005, 2:5–11) (emphasis added). Being capable of a RAID configuration is not the same as an actual RAID. To rebut Patent Owner’s arguments that Hathorn’s reference to a RAID is not enough to conclude that Hathorn discloses a RAID, Petitioner argued at the oral hearing that Figure 3, alone, shows a RAID configuration because it shows mirroring. Tr. 12:9–13:2. Petitioner alleges that Hathorn discloses that dual copy DASDs, where data is written to an additional DASD, is “sometimes referred to as mirroring.” Ex. 1005, 1:65–67. We are not persuaded by Petitioner’s argument that because DASDs are made up of disks and because a RAID is a redundant array of inexpensive disks “that DASDs can be arranged as a RAID.” See Tr. 13:12–18.

RAID is described in Hathorn as a specific architecture that is “[a]nother data back-up alternative” to DASD. Ex. 1005, 1:60–2:4, 2:4–7. The evidence, thus, supports a distinction, in the reference upon which the challenge is based, between RAID and DASD.

The record does support that RAID level 1, one of several possible RAID configurations, is “disk mirroring.” Ex. 1006, 4, 22; Ex. 1003, 153; Ex. 2007, 37:3–11. Further, Hathorn describes its “remote data shadowing” as “mirroring.” Ex. 1005, 1:65–67. However, although Hathorn may disclose a RAID level 1 configuration, the question before us is whether Hathorn discloses a

IPR2013-00635

Patent 6,978,346 B2

RAID as we have construed the term, i.e., a “single logical unit.”

Petitioner now argues that DASDs 326, 336 are a RAID. Pet. Reply 3. We agree with Patent Owner that DASDs 326, 336 are not a single logical unit because primary host 301 *can* individually (or directly) access either one of those DASDs. *See* Tr. 50:4–7.

Furthermore, Petitioner’s argument presented at oral hearing as to what is a RAID, DASDs 326, 336, differs from Petitioner’s contention presented in the Petition, where Petitioner cited *generally* to DASDs as RAIDs. Pet. 47 (annotated Fig. 3). The late change in position raises a concern that when arising so late in the proceeding the timing of presenting rebuttal evidence may be unfair and imbalanced, especially in light of our rules requiring that the Petition must specify where each element of the claim is found in the prior art patent. 37 C.F.R. §42.104(b)(4). We address, nevertheless, the merits of Petitioner’s latest position, recognizing that Patent Owner rebutted the new arguments sufficiently to dispel the prejudice that may have been sustained and would have precluded our consideration.

Based on the foregoing, we are persuaded that Petitioner has not shown by a preponderance of the evidence that Hathorn discloses a RAID as we have construed the term. Independent claim 1 recites RAID, and claims 2–3 and 5–8 all depend, directly or indirectly, from claim 1 and thus necessarily recite RAID. Petitioner has failed to show by a preponderance of the evidence that claims 1–3 and 5–8 of the ’346 patent are anticipated by Hathorn.

IPR2013-00635

Patent 6,978,346 B2

III. ORDER

For the reasons given, it is

ORDERED that claims 1–3 and 5–8 of U.S. Patent No. 6,978,346 have not been shown by a preponderance of the evidence to be unpatentable; and

FURTHER ORDERED that, because this is a final written decision, parties to the proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

IPR2013-00635
Patent 6,978,346 B2

For PETITIONERS:

Lead Counsel
David McCombs
david.mccombs.ipr@haynesboone.com

Back-up Counsel
Andrew S. Ehmke
andy.ehmke.ipr@haynesboone.com

Thomas W. Kelton
Thomas.kelton.ipr@haynesboone.com

John Russell Emerson
Russell.emerson.ipr@haynesboone.com

For PATENT OWNER:

Lead Counsel
Mathew Phillips
matthew.phillips@renaissanceiplaw.com

Back-up Counsel
Alexander Giza
agiza@raklaw.com

Derek Meeker
derek.meeker@renaissanceiplaw.com

(12) **United States Patent**
Baek et al.

(10) **Patent No.:** **US 6,978,346 B2**
 (45) **Date of Patent:** **Dec. 20, 2005**

(54) **APPARATUS FOR REDUNDANT INTERCONNECTION BETWEEN MULTIPLE HOSTS AND RAID**

6,609,213 B1 * 8/2003 Nguyen et al. 714/4
 6,820,171 B1 * 11/2004 Weber et al. 711/114

OTHER PUBLICATIONS

(75) Inventors: **Sung-Hoon Baek**, Taejon (KR);
Joong-Bae Kim, Taejon (KR);
Yong-Youn Kim, Taejon (KR)

Tellurite glass: a new candidate for fiber devices by J.S. Wang: Optical Materials 3 (1994) 187-203.
 Structure and optical properties of rare earth doped zinc oxyhalide tellurite glasses by D.L. Sidebottom: Journl of Non-Crystalline Solids 222 (1997), pp. 282-289.
 Raman spectra and thermal analysis of a new lead-tellurium-germanate glass system by Z. Pan: Journl of Non-Crystalline Solids 210 (1997), pp. 130-135.

(73) Assignee: **Electronics and Telecommunications Reseach Institute**, (KR)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 856 days.

* cited by examiner

(21) Appl. No.: **09/753,245**

Primary Examiner—Krisna Lim

(22) Filed: **Dec. 29, 2000**

(74) *Attorney, Agent, or Firm*—Blakely Sokoloff Taylor & Zafman

(65) **Prior Publication Data**

US 2002/0035669 A1 Mar. 21, 2002

(57) **ABSTRACT**

(30) **Foreign Application Priority Data**

Sep. 19, 2000 (KR) 2000-54807

The apparatus for a redundant interconnection between multiple hosts and a redundant array of inexpensive disks (hereinafter, referred to as 'RAID'), which is capable of supporting a fault tolerance of RAID controllers and simultaneously heightening a performance, comprises a plurality of RAID controlling units for processing a requirement of numerous host computers connected with one another through the industrial standard communication network and for fault tolerance; a plurality of connecting units for connecting the plurality of RAID controlling units to the numerous host computers; and a plural number of network interface controlling units respectively contained into the plurality of RAID controlling units, for exchanging information directly with an opposite network interface controlling unit provided within an opposite RAID controlling unit and the numerous host computers, through the plurality of connecting units.

(51) **Int. Cl.**⁷ **G06F 13/00; G06F 12/00**

(52) **U.S. Cl.** **711/114; 709/250**

(58) **Field of Search** 709/201-203,
 709/217-219, 223-224, 239-240, 244, 250;
 711/114; 710/38; 370/360, 412

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,652,536 A 3/1987 Nakajima et al.
 5,251,062 A 10/1993 Snitzer et al.
 5,798,306 A 8/1998 Dickinson, Jr.
 5,812,754 A * 9/1998 Lui et al. 714/6
 6,192,485 B1 * 2/2001 Takita et al. 714/6

9 Claims, 6 Drawing Sheets

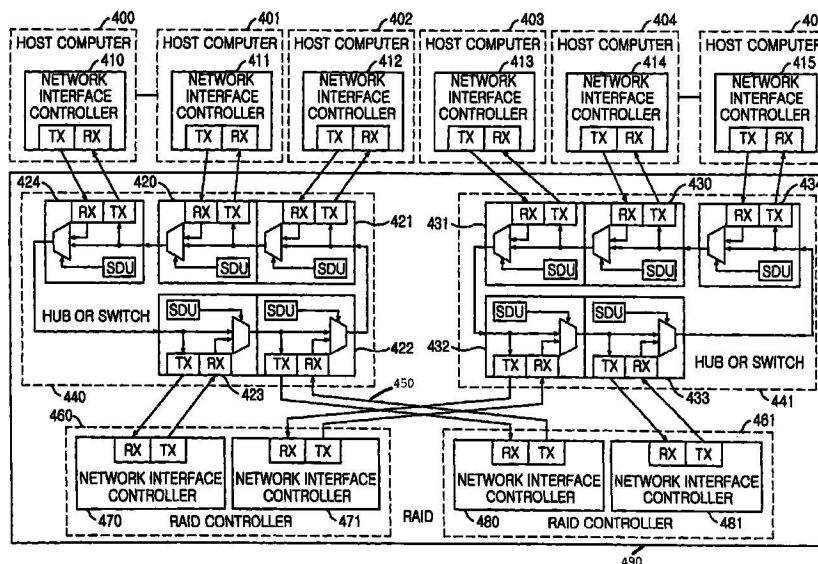


FIG. 1

--PRIOR ART--




FIG. 2

--PRIOR ART--




FIG. 3

--PRIOR ART--




FIG. 4




FIG. 5





FIG. 6



US 6,978,346 B2

1

APPARATUS FOR REDUNDANT INTERCONNECTION BETWEEN MULTIPLE HOSTS AND RAID

FIELD OF THE INVENTION

The present invention relates to an apparatus for a redundant interconnection between multiple host computers and a redundant arrays of inexpensive disks (hereinafter, referred to as 'RAID'); and, more particularly, to an apparatus for a redundant interconnection between multiple host computers and multiple controllers of the RAID, which is capable of supporting a fault tolerance of the RAID controllers and simultaneously heightening performance.

PRIOR ART OF THE INVENTION

A RAID is a storage system based on a large capacity and a high performance, by using much quantity of disks, and is a fault tolerant system in which the disks or controllers etc. have a redundant nature. In general, the RAID has two controllers, which are used like a method shown in FIG. 1 or 2.

FIG. 1 is an exemplary block diagram showing a general connection method between the host computers and the RAID having the conventional two controllers.

As shown in the drawing, the RAID 130 includes two RAID controllers 140, 141 and each of RAID controllers 140, 141 includes network interface controllers 150, 151. The network interface controllers 150, 151 of the RAID controllers 140, 141 are independently connected to network interface controllers 110, 111 of the host computers 100, 101 through communication links 120, 121 such as a copper line and an optical fiber. That is, such system has twice the bandwidth and twice the performance. However, there is such a problem that a loss of data occurs when one out of two RAID controllers 140, 141 has a trouble, in other words, this system does not become the fault tolerant system.

FIG. 2 is an exemplary block diagram of a general host interface system having a communication interface for an error recovery between the conventional two controllers.

In order to provide fault tolerance not provided in FIG. 1, RAID 240 includes two RAID controllers 230, 231 and two RAID controllers 230, 231 and host computers 200, 201 are connected with each other through a hub or switch 210 in one network. The RAID controller 230 includes a pair of network interface controllers 220 and 221 and the RAID controller 231 includes a pair of network interface controllers 222 and 223. Thus, even though one RAID controller 230 or 231 has a trouble, all of the host computers 200, 201 are connected to a RAID controller that does not have a trouble. That is, this RAID controller not having the trouble serves as a role of the controller that has the trouble. Also, since the RAID controllers 230, 231 should exchange information with each other by preparing in advance against some trouble, the RAID controllers 230, 231 are connected with each other through communication controllers 221, 222. However, in this case only a half of performance for the bandwidth provided in FIG. 1 can be obtained.

FIG. 3 is an exemplary block diagram showing a wiring method between a conventional RAID and the host computers.

The construction shown in the drawing partially represents a systematic connection between a RAID and host computers, which is extracted from contents disclosed in the U.S. Pat. No. 5,812,754. The RAID 340 includes two RAID controllers each of which has network interference control-

2

lers 330, 331 and four ports 310, 311, 320 and 321. However, this construction has no any difference from that of FIG. 2, in the structure of a communication network, and in case that one out of two host computers 300, 301 has rather a trouble, there is caused a problem that a network is broken. Thus, this construction is inferior to the construction of FIG. 2.

SUMMARY OF THE INVENTION

Therefore, it is an object of the present invention to provide an apparatus for a redundant interconnection between multiple host computers and a RAID, which is capable of supporting a fault tolerance of a RAID controller and simultaneously heightening a performance.

In accordance with the present invention, the apparatus for a redundant interconnection between multiple hosts and a RAID comprises a plurality of RAID controllers for processing requests of numerous host computers connected with one another through an industrial standard communication network such as fibre channel and performing fault tolerant function; a plurality of connecting units for connecting the plurality of RAID controllers to the numerous host computers; and a plural number of network interface controllers respectively contained into the plurality of RAID controllers, the network interface controllers being for exchanging information directly with each of opposite network interface controllers provided within the numerous host computers and within opposite RAID controllers, through the plurality of connecting units.

BRIEF DESCRIPTION OF THE DRAWINGS

The above and other objects and features of the instant invention will become apparent from the following description of preferred embodiments taken in conjunction with the accompanying drawings, in which:

FIG. 1 is an exemplary block diagram showing a general connection system between host computers and a RAID having conventional two controllers;

FIG. 2 indicates an exemplary block diagram of a general host interface system having a communication interface for an error recovery between the conventional two controllers;

FIG. 3 illustrates an exemplary block diagram of a wiring method between a conventional RAID and host computers;

FIG. 4 is a block diagram showing one embodiment of a host interface system as an internal installment system between a RAID and host computers in accordance with the present invention;

FIG. 5 depicts a block diagram providing one embodiment of a host interface system as an external installment system between a RAID and host computers in the present invention; and

FIG. 6 is a block diagram showing one embodiment of a host interface system as a network switch between a RAID and host computers in the invention.

PREFERRED EMBODIMENT OF THE INVENTION

Hereinafter, preferred embodiments of the present invention will be described in detail with reference to the accompanying drawings.

FIG. 4 is a block diagram showing one embodiment of a host matching system as an internal installment system between a RAID and host computers in accordance with the present invention.

US 6,978,346 B2

3

As shown in FIG. 4, in the inventive host interface system, a communication circuit is provided in order for an error recovery between two RAID controllers 460, 461, and the bandwidth between two groups as the host computers 400 to 405 and two RAID controllers 460, 461 becomes twice the single connection bandwidth. Also, in the inventive host interface system, even though one RAID controller 460 or 461 has an occurrence of a trouble, the bandwidth becomes twice the single connection bandwidth.

That is to say, in a RAID 490, two RAID controllers 460, 461 and hubs 440, 441 exist, and in each of the RAID controllers 460, 461, a pair of network interface controllers 470, 471; 480, 481 are provided. Herewith, the hubs 440, 441 are provided to connect a system connected to these hubs by one network and maintain the network even though one system has an occurrence of a trouble or a short of a line, and it can be as a hub or a switch. Hereinafter, they are named a "hub" altogether.

Hub ports, 420 to 424, 430 to 434, shown in FIG. 4 indicate an example for a simple internal structure of a fibre channel arbitrated loop hub, and this is based on an already well-known technique, thus there will be herein no more description therefore in the invention. The hub observes its corresponding communication network standard.

A network, in which the RAID controllers, the hubs and the host computers are connected with one another, corresponds to the industrial standard communication network such as fibre channel, asynchronous transfer mode (ATM) and InfiniBand etc. and they are hereinafter named a 'network'.

Network interface controllers, 410 to 415, contained into the host computers, 400 to 405, and the network interface controllers 470, 471, 480, 481 of the RAID controllers 460, 461 are connected with one another by two networks through two hubs 440, 441, and according to a sort of the networks, the network interface controller becomes a fibre channel controller, an ATM controller and an InfiniBand controller etc.

At this time, a communication line, representatively shown as 450 in the drawing, for connecting the network interface controller to the hub is a copper line or an optical fibre, which is matched to a corresponding standard.

Meanwhile, two network interface controllers 470, 471 of the first RAID controller 460 are respectively connected to two different hub ports 423, 432, and two network interface controllers 480, 481 of the second RAID controller 461 are respectively connected to two different hub ports 422, 433. The rest ports 420, 421, 424, 430, 431, 434 of the hubs 440, 441 are connected to the host computers 400 to 405. Just, there is no distinction between the hub ports 420 to 424 of the first hub 440 at all. Also, there is no distinction between the hub ports 430 to 434 of the second hub 441 at all.

The hub port connected to the host computer among the hub ports of the hub 440, namely, 420, 421, 424, is more than one, and there is no limitation to the maximum number. Further, What it is connected to the host computer among the hub ports of the second hub 441, namely, 430, 431, 434, is more than one, and there is no limitation to the maximum number. The hub ports 424, 434 and the host computers 400, 405, which are shown as dot lines in FIG. 4, mean that there is no, or more than one hub port or host computer.

Since, in such construction, two independent networks are constructed; it has twice the bandwidth of the single network, and a communication passage between two RAID controllers needed to perform the fault tolerant function of two RAID controllers 460, 461 is formed. Thus, information from the second network interface controller 471 of the first

4

RAID controller 460 is sent to the first network interface controller 481 of the second RAID controller 461. Also, information from the second network interface controller 480 of the second RAID controller 461 is transmitted to the first network interface controller 470 of the first RAID controller 460. Further, information from the first network interface controller 481 of the second RAID controller 461 is transmitted to the second network interface controller 471 of the first RAID controller 460, and information from the first network interface controller 470 of the first RAID controller 460 is sent to the second network interface controller 480 of the second RAID controller 461.

The first network interface controllers 470, 480 of two RAID controllers 460, 461 respectively supply data of the host computers 400 to 402 connected to the first hub 440 and the host computer 403 to 405 connected to the second hub 441, and process information transmitted from the opposite network interface controllers 471, 481.

If any one out of two RAID controllers 460, 461 has an occurrence of an error, the RAID controller having the error occurrence is removed from the network, and a second network interface controller of an opposite RAID controller not having the error occurrence takes over a function of a first network interface controller of the RAID controller having the error occurrence.

FIG. 5 is a block diagram providing one embodiment of the host interface system as an external installation system between the RAID and the host computers in the present invention.

As shown in FIG. 4, the present invention can be constructed by a method of internally installing the hubs 440, 441 in the RAID 490, and as shown in FIG. 5, the host computers 500, 501, 502, 503, 504 and 505 are connected to the RAID 530 by using external hubs 510 and 520.

FIG. 6 is a block diagram showing one embodiment of the host interface system as a network switch between the inventive RAID and host computers.

As shown in the drawing, a plurality of host computers 600, 601, 602, 604 and 605 are connected to RAID through a network switch 610. In other words, information from a second network interface controller 622 of a first RAID controller 620 is sent to a first network interface controller 632 of a second RAID controller 630, and information from a second network interface controller 632 of the second RAID controller 630 is transmitted to a first network interface controller 621 of the first RAID controller 620. Further, information from the first network interface controller 631 of the second RAID controller 630 is transmitted to the second network interface controller 622 of the first RAID controller 620. Also, information from the first network interface controller 621 of the first RAID controller 620 is sent to the second network interface controller 632 of the second RAID controller 630.

Just, there is no distinction between respective ports, representatively 611, of a network switch 610 at all and also, the internal structure of a network switch 610 can be configured according to a selection of a user (not shown in FIG. 6).

In accordance with the present invention, as aforementioned, even in a case of an error occurrence in a RAID controller, there exist two independent networks and two network interface controllers, and the bandwidth of a single network can be twice maintained. Accordingly, a function of fault tolerance between two RAID controllers can be constructed without a drop of the bandwidth.

It will be apparent to those skilled in the art that various modifications and variations can be made in the present

US 6,978,346 B2

5

invention without deviating from the spirit or scope of the invention. Thus, it is intended that the present invention cover the modifications and variations of this invention provided they come within the scope of the appended claims and their equivalents.

What is claimed is:

1. An apparatus for a redundant interconnection between multiple hosts and a RAID, comprising:

a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers, the first RAID controlling unit including a first network controlling unit and a second network controlling unit, and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit; and

a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers, wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the numerous host computers through the plurality of connecting units, and the first network controlling unit exchanges information with the fourth network controlling unit, and the second network controlling unit exchanges information with the third network controlling unit.

2. The apparatus as recited in claim 1, wherein said respective RAID controlling units are connected to the plurality of individual connecting units.

3. The apparatus as recited in claim 2, wherein the first network interface controlling unit is coupled to the connecting unit of one side and the second network interface controlling unit is coupled to the connecting unit of another side.

4. The apparatus as recited in claim 3, wherein the first network interface controlling unit and the third network interface controlling unit process the requirement of the numerous host computers; and the second network interface controlling unit and the fourth network controlling unit are used for communication between the first RAID controlling unit and the second RAID controlling unit when the first and second RAID controlling units are not faulty and the second network interface controlling unit and the fourth network controlling unit are used for executing a function of the first network interface controlling unit and the third network controlling unit when one of the first RAID controlling unit and the second RAID controlling unit is faulty.

5. The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports, two of the at least three connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit and the rest of the connection ports being provided as a hub equipment connected with the numerous host computers.

6. The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports, two of the at least three connection port are coupled to one of the first network interface controlling unit and the third network controlling unit and the rest of the connection

6

ports being provided as a network switch equipment connected with the numerous host computers.

7. The apparatus as recited in claim 1, wherein said plurality of connecting units have at least five connection ports, four of the at least five connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit and the rest of the connection ports being provided as a switch connected with the numerous host computers.

8. The apparatus as recited in claim 1, wherein the first network interface controlling unit of the first RAID controlling unit being connected to a first connecting unit, the second network interface controlling unit of said first RAID controlling unit being connected to a second connecting unit, the third network interface controlling unit of the second RAID controlling unit being connected to the second connecting unit, and the fourth network interface controlling unit of the second RAID controlling unit being connected to the first connecting unit.

9. An apparatus for a redundant interconnection between multiple host computers and a RAID, the apparatus comprising:

a plurality of connection units for connecting the host computers and the RAID;

a first and a second RAID controllers, included in the RAID, each of which having a first network interface controller and a second network interface controller for processing requests from the plurality of the host computers connected through the plurality of the connection units,

wherein the first network interface controller in the first RAID controller supplies data to the host computers connected through the plurality of connection units and processes information transmitted from the second network interface controller in the second RAID controller,

wherein the first network interface controller in the second RAID controller supplies data to the host computers connected through the plurality of connection units and processes information transmitted from the second network interface controller in the first RAID controller,

wherein the second network interface controller in the first RAID controller is used for fault tolerance by performing functions of the first network interface controller in the second RAID controller when the second RAID controller is faulty, and

wherein the second network interface controller in the second RAID controller is used for fault tolerance by performing functions of the first network interface controller in the first RAID controller when the first RAID controller is faulty, and

wherein the first network controlling unit in the first RAID controlling unit exchanges information with the second network controlling unit in the second RAID controlling unit, and the second network controlling unit in the first RAID controlling unit exchanges information with the first network controlling unit in the second RAID controlling unit.

* * * * *